Discrete Load Balancing in Heterogeneous Networks with a Focus on Second-Order Diffusion

Hoda Akbari* Petra Berenbrink*
*Simon Fraser University
{hodaa, petra}@sfu.ca

Robert Elsässer† Dominik Kaaser†
†University of Salzburg
{elsa, dominik}@cosy.sbg.ac.at

arXiv:1412.7018v1 [cs.DC] 22 Dec 2014

***Abstract*—In this paper we consider a wide class of discrete diffusion load balancing algorithms. The problem is defined as follows. We are given an interconnection network and a number of load items, which are arbitrarily distributed among the nodes of the network. The goal is to redistribute the load in iterative discrete steps such that at the end each node has (almost) the same number of items. In diffusion load balancing nodes are only allowed to balance their load with their direct neighbors.**

We show three main results. Firstly, we present a general framework for randomly rounding the flow generated by continuous diffusion schemes over the edges of a graph in order to obtain corresponding discrete schemes. Compared to the results of Rabani, Sinclair, and Wanka, FOCS'98, which are only valid w.r.t. the class of homogeneous first order schemes, our framework can be used to analyze a larger class of diffusion algorithms, such as algorithms for heterogeneous networks and second order schemes. Secondly, we bound the deviation between randomized second order schemes and their continuous counterparts. Finally, we provide a bound for the minimum initial load in a network that is sufficient to prevent the occurrence of negative load at a node during the execution of second order diffusion schemes.

Our theoretical results are complemented with extensive simulations on different graph classes. We show empirically that second order schemes, which are usually much faster than first order schemes, will not balance the load completely on a number of networks within reasonable time. However, the maximum load difference at the end seems to be bounded by a constant value, which can be further decreased if first order scheme is applied once this value is achieved by second order scheme.

I. Introduction

Load balancing is a fundamental task in many parallel and distributed applications. Often there are significant differences in the amount of work load generated on the processors of a parallel machine, which have to be balanced in order to obtain a substantial benefit w.r.t. the runtime of a parallel computation. One of the most prominent examples are so-called finite element simulations [14].

In the load balancing problem we are given an interconnection network and a number of load items which are arbitrarily distributed over the nodes of the network. The goal is to redistribute the items such that at the end each node has (almost) the same load. To achieve this goal, nodes are only allowed to communicate with their direct neighbors. We assume that each node has access to a global clock, and the algorithm works in synchronous rounds.

A prominent class of load balancing algorithms are so-called diffusion schemes [8]. In these algorithms, the nodes are allowed to balance their load with all their neighbors simultaneously in a round. We distinguish between *continuous* and *discrete* settings. In the continuous case it is assumed that the load can be split into arbitrarily small pieces. Although often not realistic, this assumption is very helpful for analyzing these algorithms [8]. Discrete load balancing algorithms, on the other hand, assume that tasks are atomic units of load, called *tokens*. Hence, two adjacent nodes cannot balance their load any way they want; only integral amounts of load can be transferred. As a consequence, discrete diffusion algorithms are usually not able to balance the load completely [2], [12].

Two fundamental diffusion type algorithms are the first order scheme (FOS) and the second order scheme (SOS) [19]. In the first order scheme the amount of load that nodes send to their neighbors in a step only depends on their current load difference. In SOS the flow over an edge is a function of the current load difference between its incident nodes and the load that was sent in the previous round. Note that SOS can lead to negative load at some nodes if the loads of the nodes are not sufficient to fulfill the calculated demand of all edges. There are tight bounds on the worst-case convergence time of both, FOS and SOS, in the continuous case [8]. In general, for the optimal choice of parameters SOS converges much faster than FOS.

The common approach for analyzing discrete diffusion algorithms is to consider a closely related continuous version of the algorithm and to bound the load deviation between load vectors of the two processes ([20]). To explain the approach we need a couple of definitions first. We assume that the network is modeled by an undirected graph $G = (V, E)$, where $V = \{1, \ldots, n\}$ represents the set of processors and the edges in E describe the connections between them. A total of m identical load items are distributed over the nodes. We use a vector $x = (x_1, \ldots, x_n)$ to indicate the amount of load assigned to every node. In the *heterogeneous* network model the nodes may have different speeds $(s_1, \ldots, s_n)$. The aim of a load balancing algorithms is to distribute the load proportional to the processors' speeds. Hence, the ideal load of a node i is $\bar{x}_i = ms_i/s$, where $s = \sum_{i=1}^{n} s_i$. The deviation of a load vector x from another load vector x' is $\max_{i \le n} |x_i - x'_i|$.

In the case of the common approach mentioned above the continuous process would forward a fractional amount of load ℓ_e over some edge e, the discrete algorithm rounds ℓ_e to an integer ℓ'_e. The rounding can be done deterministically or randomized, whereas randomized rounding often outperforms deterministic rounding (for example, the always round down approach [21]). The difference between ℓ_e and ℓ'_e is called the

rounding error. The propagation of the rounding errors causes the two processes to deviate from each other.

In this paper we show three main results. Firstly, we present a general framework for randomly rounding continuous diffusion schemes to discrete schemes. Compared to the results of [20], which are only valid w.r.t. the class of homogeneous first order schemes, our framework can be used to analyze a larger class of diffusion algorithms, such as algorithms for heterogeneous networks and second order schemes. Secondly, we bound the deviation between randomized second order schemes and their continuous counterparts. Finally, we provide a bound for the minimum initial load in a network that is sufficient to prevent the occurrence of negative load at a node during the execution of second order diffusion schemes. Our results are supported by extensive simulations on various graph classes, comparing the performance of FOS and SOS and giving an empiric insight into the behavior of diffusion based load balancing processes.

II. Models and Results

a) First Order Diffusion: FOS in the *homogeneous* network model is defined as follows. Let $N(i)$ be the set of neighbors of node i and d_i be its degree. We define $x(t) = (x_1(t), \ldots, x_n(t))$ as the *load vector* at the beginning of round $t \geq 0$, where $x_i(t)$ is the load of node i. The amount of load transferred from node i to node j in round t is denoted by $y_{i,j}(t)$. Then FOS is characterized by the following equations, where $\alpha_{i,j}$ is a parameter, usually $\alpha_{i,j} = 1/\left(\max(d_i, d_j) + 1\right)$.

$$\begin{aligned} y_{i,j}(t) &= \alpha_{i,j} \cdot \big(x_i(t) - x_j(t)\big) \\ x_i(t+1) &= x_i(t) - \sum_{j \in N(i)} \alpha_{i,j} \left(x_i(t) - x_j(t)\right) \end{aligned} \tag{1}$$

The process can be expressed with a *diffusion matrix* M, where $M_{i,i} = 1 - \sum_j \alpha_{i,j}$ and $M_{i,j} = \alpha_{i,j}$ for $j \in N(i)$. All other entries of M are zero. Then

$$x(t+1) = M \cdot x(t) \ , \tag{2}$$

where M is a symmetric doubly stochastic $n \times n$ matrix. Let K denote the difference between the maximum and minimum load at the beginning of the process. Let λ denote the second-largest eigenvalue (in magnitude) of M. Then [19], [20] show that FOS converges in $O\left(\log(Kn)/(1-\lambda)\right)$ rounds. In [20] the authors introduce a framework to analyze a wide class of discrete FOS processes. This framework served as a foundation for analyzing several discrete FOS algorithms. Many of these publications consider uniform processors [5], [6], [10], [13], [15], [16], [17], [19], [20], [21], while a few others incorporate processor speeds into the model [1], [12]. The authors of [15] consider a discrete process where the continuous flow is rounded randomly. This algorithm achieves a deviation bound of $O((d \log\log n)/(1-\lambda))$. The drawback of this method is that rounding up on too many edges might result in negative load. The process of [5] avoids negative load. A node first rounds down all the flows on the adjacent edges, which leaves some surplus tokens which are randomly distributed among the neighbors. This algorithm achieves a deviation bound of $O(d\sqrt{\log n} + \sqrt{(d \log n \log d)/(1-\lambda)})$. In [21] the authors study two natural discrete diffusion-based protocols and their discrepancy bounds depend only polynomially on the maximum degree of the graph and logarithmically on n.

The balancing process of [2] simulates a continuous process using a corresponding discrete process. In every round the discrete flow on each edge is determined such that it stays as close as possible to the total continuous flow that is sent over the edge. This process results in a deviation of $O(d)$ (for a more detailed description see next section). In [13] the authors consider an approach that is based on random walks where tokens of overloaded nodes use a random walk to reach underloaded nodes. While this approach leads to a situation at the end, in which no node has more than a constant number of tokens above average [13], it needs to keep track of the load traffic the continuous scheme would produce. Moreover, the corresponding random walks of the tokens result in a huge amount of load transmissions between the nodes, which is not the case in diffusion based schemes [8].

b) Second Order Diffusion (SOS): Muthukrishnan et al. [19] introduce the continuous second order scheme which is based on a numerical iterative method called successive over-relaxation [18] and is one of the fastest diffusion load balancing algorithms. In SOS, the amount of load transmitted over each edge depends on the current load as well as the load transferred in the previous round. The only exception is the very first round in which FOS is applied. Subsequent rounds follow the equations below.

$$y_{i,j}(t) = (\beta - 1)\ y_{i,j}(t-1) + \beta \alpha_{i,j} \left(x_i(t) - x_j(t)\right) \tag{3}$$

$$\begin{aligned} x_i(t+1) = \beta \cdot \left(x_i(t) - \sum_{j \in N(i)} \alpha_{i,j} \left(x_i(t) - x_j(t)\right) \right) \\ + (1-\beta) \cdot x_i(t-1) \end{aligned}$$

Here, β is independent of the iteration number t. From the above equations we get

$$x(t+1) = \begin{cases} M\, x(t) & \text{if } t = 0 \\ \beta \cdot M\, x(t) + (1-\beta) \cdot x(t-1) & \text{if } t > 0 \end{cases} \tag{4}$$

For the process to converge, β must be in the interval $(0, 2)$. For the optimal choice of $\beta_{opt} = 2/(1+\sqrt{1-\lambda^2})$ SOS converges in $O(\log(Kn)/\sqrt{1-\lambda})$ rounds [19] which is in general faster than FOS; for graphs with some eigenvalue gap $(1-\lambda)^{-1} = \log^{\omega(1)} n$, the convergence time of SOS is almost quadratically faster than FOS. Unfortunately, it can happen that the total outgoing flow from a node exceeds its current load, which results in so-called *negative load*.

c) Heterogeneous Networks: Continuous FOS and SOS processes in the *heterogeneous network model* were first studied in [11]. In heterogeneous networks, processors have different speeds and the aim is to distribute the load proportional to their speeds. The minimum speed is 1, the maximum speed is $s_{\max}$, and $s = s_1 + \cdots + s_n$. Let the diagonal matrix S be defined by $S_{i,i} = s_i$. Then the heterogeneous FOS/SOS processes are defined as before (see (2) and (4)), except the diffusion matrix is now $M = I - LS^{-1}$ where L is the normalized Laplacian matrix of the graph [11]. In [1], the authors analyze a discrete FOS for homogeneous networks. In [11] the authors show that continuous FOS/SOS processes converge in $O(\log(Kns_{\max})/(1-\lambda))$ and $O(\log(Kns_{\max})/\sqrt{1-\lambda})$

rounds, respectively. In [12], the authors consider a discrete version of SOS too. They show that the euclidean distance between the discrete and continuous load vectors in the discrete version is $O(d \cdot \sqrt{n \cdot s_{\max}}/(1-\lambda))$.

A. New Results

Result I. We present a general framework for rounding continuous diffusion schemes to discrete schemes. Our approach described in Section III estimates the error between a continuous diffusion scheme and the rounded discrete version first, similar to [20]. Then we combine that error term with martingales techniques (similar to the ones used in [5]) to bound the deviation between the continuous scheme and a discrete scheme based on randomized rounding. Note that the results in [20] are only valid for a class of homogeneous first order schemes and [5] analyzes a fixed first order diffusion scheme with a specific transition matrix. In this paper we introduce an error estimation that allows us to show results for a larger class of diffusion algorithms (see Definition 2) in heterogeneous networks, including SOS.

In the homogeneous case our bounds are the same as the best results for FOS. Our bound is worse than the $O(d)$ bound of [2]. In the current paper we bound the deviation of a class of very natural and stateless algorithms. That is, the amount of load that is forwarded over an edge in step t only depends on the load at the beginning of step t and the amount that was sent in step $t-1$. The approach of [2] is not stateless as it simulates the continuous process. The flow that is sent over the edges in step t takes into account the difference of the cumulative load that was sent by the continuous process up to step t and the cumulative load that was sent by the discrete process so far.

Result II. We show that randomized SOS has a deviation (after the balancing time of continuous SOS) of $O\left(d \cdot \log s_{\max} \cdot \sqrt{\log n}/(1-\lambda)^{3/4}\right)$, where λ is the second largest eigenvalue of M and $s_{\max}$ is the maximum speed. Note that the runtime of SOS is in most cases much better than the runtime of FOS, i.e., $O(\log(Kn)/\sqrt{1-\lambda})$ (assuming optimal β) compared to $O\left(\log(Kn)/(1-\lambda)\right)$ in the case of FOS.

Result III. We show that the continuous second order scheme with optimal β will not generate negative load if at time $t=0$ the minimum load of every node is at least $O\left(\sqrt{n} \cdot \Delta(0)/\sqrt{1-\lambda}\right)$. Here $\Delta(0)$ is the difference between the maximum load and the average load at time $t=0$. For discrete SOS and graphs with proper eigenvalue gap we show a bound of $O\left((\sqrt{n} \cdot \Delta(0) + d^2)/\sqrt{1-\lambda}\right)$. To the best of our knowledge these are the first results specifying a sufficient amount of minimum load w.r.t. SOS to avoid negative load.

Simulations. We implemented a network and simulated both, FOS and SOS load balancing processes. Especially in tori, our results show a clear advantage of SOS over FOS w.r.t. the number of steps required to balance the loads. We also empirically analyze the remaining imbalance that arises in discrete load balancing schemes once the system has converged such that no node has more than a constant number of additional load tokens. We propose to switch from SOS to FOS once this threshold is reached, and our simulations show that this change of the scheme leads to a further drop of the remaining load imbalance.

III. General Framework for FOS Schemes

In this section we first generalize the framework of Rabani et al. [20] to a wider class of processes (see Section III-A) and obtain an equation estimating the deviation of the discrete process from its continuous version. The estimation is valid as long as the continuous process is *linear* (Definition 2). In [20] the deviation is expressed in terms of the diffusion matrix. Here, we present an analysis from a different perspective which allows us to obtain essentially the same deviation formula for a larger class of processes. Our analysis can be applied to the second order processes and heterogeneous models. In Section III-B we present the framework that transforms a continuous load balancing process C into a discrete process $R(C)$ using randomized rounding.

For simplicity we consider in this section only first order processes. In the next section we generalize the framework to SOS.

A. Deviation between Continuous and Discrete FOS

For a load balancing process A, we use $x_i^A(t)$ to denote the load of a node i at the beginning of the round t, and $x^A(t) = (x_1^A(t), \ldots, x_n^A(t))$. For $j \in N(i)$, we define $y_{i,j}^A(t)$ as the amount of load sent from i to j in round t (this value is negative if load items are transferred from j to i), where $N(i)$ represents the set of neighbors of i. Then $y^A(t)$ is the matrix with $y_{i,j}^A(t)$ as its entry in row i and column j. Note that each balancing process A can be regarded as a function that, given the current state of the network, determines for every edge e and round t the amount of load that has to be transferred over e in t. Hence, we can regard $y^A(t)$ as the result of applying a function A, i.e., $y^A(t) = A(x^A(t))$. Using this we formally define discrete processes as follows.

Definition 1. *Let C be a continuous process. A process D is said to be a discrete version of C with rounding scheme R_D if for every vector $\mathbf{x}$, we have $D(\mathbf{x}) = R_D(C(\mathbf{x}))$ where R_D is a function that rounds each entry of the matrix to an integer.*

Note that for a load balancing process load conservation over each edge must hold. Although it may not be a necessary condition, our analyses in this section requires the process to exhibit a *linearity* property in the following sense.

Definition 2 (Linearity)**.** *A diffusion process A is said to be* linear *if for all $\mathbf{x}, \mathbf{x}' \in \mathbb{R}^n$ and $a, b \in \mathbb{R}$ we have $A(a\mathbf{x}+b\mathbf{x}') = a \cdot A(\mathbf{x}) + b \cdot A(\mathbf{x}')$.*

Lemma 1. *Both FOS and SOS as defined in Section II are linear.*

Proof: Let M be the diffusion matrix and $0 \leq \beta \leq 2$. Observe that both FOS and SOS can be described by the following general equation (see equations (1) and (3)).

$$y_{i,j}(t) = (\beta-1) \cdot y_{i,j}(t-1) + \beta \cdot M_{i,j} \cdot x_i(t) \quad \text{for } t \geq 1,$$

Thus the algorithm A – where based on the choice of parameter β, A can represent either FOS and SOS – is defined by

$$A(\mathbf{x}, \mathbf{y}) = (\beta-1)\mathbf{y} + \beta M \mathbf{x}$$

Let $\mathbf{x}, \mathbf{x}' \in \mathbb{R}^n, \mathbf{y}, \mathbf{y}' \in \mathbb{R}^{n\times n}$ and $a, b \in \mathbb{R}$. Then we have

$$\begin{aligned} A(a\mathbf{x} + b\mathbf{x}', a\mathbf{y} + b\mathbf{y}') &= (\beta - 1)(a\mathbf{y} + b\mathbf{y}') + \beta M(a\mathbf{x} + b\mathbf{x}') \\ &= a((\beta-1)\mathbf{y} + \beta M\mathbf{x}) + b((\beta-1)\mathbf{y}' + \beta M\mathbf{x}') \\ &= aA(\mathbf{x}, \mathbf{y}) + bA(\mathbf{x}', \mathbf{y}') \end{aligned}$$

which shows that A is linear. ∎

Let C be a continuous process and D its discrete version. Let $\hat{Y}(t)$ represent $C(x^D(t))$. Then we can say that D always *attempts to* set $y^D_{i,j}(t)$ to $\hat{Y}_{i,j}(t)$. Hence, we call $\hat{Y}(t)$ the *continuous scheduled load*. We define the *rounding error* as $e_{i,j}(t) = \hat{Y}_{i,j}(t) - y^D_{i,j}(t)$; note that $e_{i,j}(t) = -e_{j,i}(t)$.

In the next definition $\hat{\mathbf{i}}$ denotes the unit vector of length n with 1 as its i'th entry.

Definition 3 (Contributions)**.** *Let* $\mathbf{x}$ *and* $\mathbf{x}'$ *be the load vectors obtained from applying* C *for* t *rounds on* $\hat{\mathbf{i}}$ *and* $\hat{\mathbf{j}}$*, respectively. For two fixed nodes* i *and* k *and* $j \in N(i)$ *the* contribution *of edge* (i, j) *on node* k *after* t *rounds is defined as*

$$\mathcal{C}^C_{k,i\to j}(t) = \mathbf{x}_k - \mathbf{x}'_k \ .$$

The next theorem provides a general form of the FOS deviation formula of [20] which has served as a basis for analyzing several discrete FOS processes.

Lemma 2. *Consider a linear diffusion process* C *and its discrete version* D *with an arbitrary rounding scheme. Then, for an arbitrary node* k *and round* t *we have*

$$x^D_k(t) - x^C_k(t) = \sum_{s=1}^{t} \sum_{\{i,j\}\in E} e_{i,j}(t-s)\, \mathcal{C}^C_{k,i\to j}(s)$$

Proof: Fix a node k and round t. Suppose we *sequentialize* the load balancing actions of the process by imposing an arbitrary ordering on the edges. Then, t rounds in the parallel view is equivalent to $|E| \cdot t$ steps in the sequentialized view. In the following, let $\tau = |E| \cdot t$. With a slight abuse of notation we let $C^\infty \circ D^\ell$ denote a *hybrid* process in which the load balancing actions are determined by D in steps 1 to ℓ and by C afterwards, where $0 \le \ell \le \tau$. Observe that $x^{C^\infty\circ D^\tau}_k(t) = x^D_k(t)$ and $x^{C^\infty\circ D^0}_k(t) = x^C_k(t)$. Thus we can write $x^D_k(t) - x^C_k(t)$ in the form of a telescoping sum as follows.

$$\begin{aligned} x^D_k(t) - x^C_k(t) &= x^{C^\infty\circ D^\tau}_k(t) - x^{C^\infty\circ D^0}_k(t) \\ &= \sum_{\ell=1}^{\tau} \left(x^{C^\infty\circ D^\ell}_k(t) - x^{C^\infty\circ D^{\ell-1}}_k(t) \right) \end{aligned} \quad (5)$$

Fix an arbitrary step ℓ and let $\{i, j\}$ and s be the edge and the round corresponding to the step ℓ. Both $C^\infty \circ D^{\ell-1}$ and $C^\infty \circ D^\ell$ start their round $s + 1$ with load vectors that are the same except maybe in i and j. This happens because $C^\infty \circ D^{\ell-1}$ forwards $\hat{Y}_{i,j}(s)$ over $\{i, j\}$ while in $C^\infty \circ D^\ell$ this amount is $\hat{Y}_{i,j}(s) - e_{i,j}(s)$. Thus, by the definition of $\mathcal{C}^C_{k,i\to j}(t)$ and using the linearity property of the process we get

$$x^{C^\infty\circ D^\ell}_k(t) - x^{C^\infty\circ D^{\ell-1}}_k(t) = e_{i,j}(s)\, \mathcal{C}^C_{k,i\to j}(t-s) \ .$$

Plugging the above into (5) and translating the summation index we get

$$\begin{aligned} x^D_k(t) - x^C_k(t) &= \sum_{s=0}^{t-1} \sum_{\{i,j\}\in E} e_{i,j}(s)\, \mathcal{C}^C_{k,i\to j}(t-s) \\ &= \sum_{s=1}^{t} \sum_{\{i,j\}\in E} e_{i,j}(t-s)\, \mathcal{C}^C_{k,i\to j}(s) \ . \end{aligned}$$

∎

B. Framework for Randomized FOS

In this section we use Lemma 2 to analyze a randomized rounding scheme for a general class of continuous load balancing algorithms. Our technique is based on the results in [5] where the authors analyzed a fixed discrete FOS process for homogeneous d-regular graphs using randomized rounding. Their algorithm is based on a continuous process in which every node sends a $1/(d + 1)$-fraction of its load to each neighbor. Initially, the discrete algorithm rounds $x_i/(d + 1)$ down if it is not an integer. This leaves $(d+1) \cdot \lfloor x_i/(d+1) \rfloor$ surplus tokens on node i, which they call *excess* tokens. The excess tokens are then distributed by sending the tokens to neighbors which are uniformly sampled without replacement.

Here we apply the technique in a much more general way, using Lemma 2 to express the deviation between the randomized and deterministic algorithm. We introduce a randomized framework that converts a general class of continuous processes to their discrete versions using randomized rounding. For $a \in \mathbb{R}$ we use $\{a\}$ to denote $a - \lfloor a \rfloor$.

The Randomized Rounding Algorithm. Fix a node i. Let $\hat{Y}(t)$ represent $C(x^D(t))$. For each edge $e = \{i, j\}$ let the corresponding $\hat{Y}_{i,j}(t)$ be the load that would be sent over e by the continuous process C. The rounding scheme works as follows. First, it rounds $\hat{Y}_{i,j}(t)$ down for all the edges. This leaves $r = \sum_{j:\hat{Y}_{i,j}(t)\ge 0} \{\hat{Y}_{i,j}(t)\}$ excess load on node i. Then it takes $\lceil r \rceil$ additional tokens and sends each of them out with a probability of $r/\lceil r \rceil$. With the remaining probability the excess tokens remain on node i. The tokens which do not remain on i are sent to a neighbor j with a probability of $\{\hat{Y}_{i,j}(t)\}/r$. Let $Z_{i,j}(t)$ be a counting random variable denoting the number of excess tokens that i sends to j in round t. Then we have

$$Y^R_{i,j}(t) = \begin{cases} \lfloor \hat{Y}_{i,j}(t) \rfloor + Z_{i,j}(t) & \text{if } \hat{Y}_{i,j}(t) \ge 0 \\ -Y^R_{j,i}(t) & \text{otherwise.} \end{cases}$$

The deviation bound is expressed based on the *refined local divergence* $\Upsilon^C(G)$ defined below, which is a function of both the algorithm and the graph:

$$\Upsilon^C(G) = \max_{k\in V} \left(\sum_{s=0}^{\infty} \sum_{i=1}^{n} \max_{j\in N(i)} \left(\mathcal{C}^C_{k,i\to j}(s) \right)^2 \right)^{1/2}$$

$\Upsilon^C(G)$ is a generalization of the refined local divergence $\Upsilon(G)$ introduced in [5]. Then we have the following result.

Theorem 3. *Let* C *be a continuous FOS and let* $R = R(C)$ *be a discrete FOS using our randomized rounding transformation.*

In an arbitrary round t we have w.h.p.[1]

$$\left|X_k^R(t) - x_k^C(t)\right| = O\left(\Upsilon^C(G) \cdot \sqrt{d \log n}\right) \ .$$

The proof of Theorem 3 relies on the fact that FOS is a linear process and hence the estimation of Lemma 2 can be used as a basis for the randomized analysis. The proof is similar to the proof of [5], the difference is that we use $\mathcal{C}^C_{k,i\to j}(t)$'s instead of the diffusion matrix. We begin the proof of Theorem 3 with a simple observation.

Observation 1. *The following statements are true (Recall that $\{a\}$ denotes $a - \lfloor a \rfloor$).*

1) If $\hat{Y}_{i,j}(t) \geq 0$ then $E_{i,j}(t) = \{\hat{Y}_{i,j}(t)\} - Z_{i,j}(t)$;

2) $\mathrm{E}\left[E_{i,j}(t)\right] = 0$.

The first statement of Observation 1 holds by definition, since $E_{i,j}(t) = \hat{Y}_{i,j}(t) - Y^R_{i,j}(t)$ while $Y^R_{i,j}(t) = \left\lfloor \hat{Y}_{i,j}(t) \right\rfloor + Z_{i,j}(t)$. For the second statement, first suppose that $\hat{Y}_{i,j}(t) \geq 0$. note that $Z_{i,j}(t)$ can be expressed as a sum of $\lceil r \rceil$ identically distributed Bernoulli random variables each of which is one with probability $(r/\lceil r \rceil) \cdot \left(\{\hat{Y}_{i,j}(t)\}/r\right)$. Thus we have $\mathrm{E}\left[Z_{i,j}(t)\right] = \{\hat{Y}_{i,j}(t)\}$ and from there (2) follows from (1). In the case $\hat{Y}_{i,j}(t) < 0$, we have $\hat{Y}_{j,i}(t) = -\hat{Y}_{i,j}(t) > 0$. Thus by the first case we have $\mathrm{E}\left[E_{j,i}(t)\right] = 0$ and therefore $\mathrm{E}\left[E_{i,j}(t)\right] = -\mathrm{E}\left[E_{j,i}(t)\right] = 0$.

Let f_k denote the difference in the load of k in round t of R and C. In the following, we first observe that f_k is zero in expectation, and then show that it is well concentrated around its average.

Observation 2. $\mathrm{E}\left[f_k\right] = 0$

Proof: The statement follows from Lemma 2 and Observation 1.(2) by the linearity of expectation. ∎

Proof of Theorem 3: As in [5], we are going to use the method of averaged bounded differences to obtain concentration results for the random variable f_k. For a fixed initial load vector $X(0)$ the function f_k depends only on the randomly chosen destinations of the excess tokens. There are t rounds, n nodes, and at most d excess tokens per node per round. Similar to [5] we describe these random choices by a sequence of tnd random variables, $Y_1, Y_2, \ldots, Y_{tnd}$. For any ℓ with $1 \leq \ell \leq tnd$, let (s, i, b) be such that $\ell = snd + (i-1)\,d + b$ (note that (s, i, b) is the ℓ-th largest element in the sequence). Then Y_ℓ refers to the destination of the b-th excess token of vertex i in round s (if there is one). More precisely,

$$Y_\ell = \begin{cases} j & \text{if the } b\text{-th excess token of the vertex } i \text{ in} \\ & \text{round } s \text{ is sent to } j, \text{ and} \\ 0 & \text{otherwise.} \end{cases}$$

Let $\mathbb{Y}_{\boldsymbol{i}}$ denote $Y_i, \ldots, Y_1$. To apply the method of averaged bounded differences, we need to bound the difference sequence below.

$$\left|\mathrm{E}\left[f_k \mid \mathbb{Y}_{\boldsymbol{\ell}}\right] - \mathrm{E}\left[f_k \mid \mathbb{Y}_{\boldsymbol{\ell-1}}\right]\right|. \tag{6}$$

[1]Throughout this paper, w.h.p. means with probability at least $1 - n^{-\alpha}$ for some constant $\alpha > 0$.

As in [5], we consider a fixed ℓ that corresponds to (s_1, i_1, b_1) in the lexicographic ordering.

To bound (6), we write

$$\begin{aligned} c_\ell &= \left|\mathrm{E}\left[f_k \mid \mathbb{Y}_{\boldsymbol{\ell}}\right] - \mathrm{E}\left[f_k \mid \mathbb{Y}_{\boldsymbol{\ell-1}}\right]\right| \\ &\leq \sum_{s=0}^{t} \sum_{\{i,j\}\in E} \left|\mathrm{E}\left[E_{i,j}(s) \mid \mathbb{Y}_{\boldsymbol{\ell}}\right] - \mathrm{E}\left[E_{i,j}(s) \mid \mathbb{Y}_{\boldsymbol{\ell-1}}\right]\right| \\ &\quad \cdot \left|\mathcal{C}^C_{k,i\to j}(t-s)\right| \end{aligned}$$

As in [5] we split the sum over s into the three parts $1 \leq s < s_1$, $s = s_1$, and $s_1 < s \leq t$. In the following we show that the sums over $s < s_1$ and $s > s_1$ are both zero while the part $s = s_1$ is upper bounded by $2 \cdot \max_{j\in N(i_1)} \left|\mathcal{C}^C_{k,i\to j}(t-s)\right|$.

Case $\mathbf{s < s_1}$: For every $\{i,j\} \in E$, $E_{i,j}(s)$ is already determined by $\mathbb{Y}_{\boldsymbol{\ell-1}}$. Hence,

$$\begin{aligned} &\sum_{s=1}^{s_1-1} \sum_{\{i,j\}\in E} \left|\mathrm{E}\left[E_{i,j}(s) \mid \mathbb{Y}_{\boldsymbol{\ell}}\right] - \mathrm{E}\left[E_{i,j}(s) \mid \mathbb{Y}_{\boldsymbol{\ell-1}}\right]\right| \\ &\cdot \left|\mathcal{C}^C_{k,i\to j}(t-s)\right| = 0 \ . \end{aligned} \tag{7}$$

Case $\mathbf{s = s_1}$: In this case, $\mathbb{Y}_{\boldsymbol{\ell-1}}$ determines $\hat{Y}_{i,j}(t)$ and $E_{i,j}(s)$ is only affected by $Z_{i,j}(s)$'s (see Observation 1).

$$\begin{aligned} &\sum_{\{i,j\}\in E} \left|\mathrm{E}\left[E_{i,j}(s) \,\middle|\, \mathbb{Y}_{\boldsymbol{\ell}}\right] - \mathrm{E}\left[E_{i,j}(s) \mid \mathbb{Y}_{\boldsymbol{\ell-1}}\right]\right| \\ &\quad \cdot \left|\mathcal{C}^C_{k,i\to j}(t-s)\right| \\ &= \sum_{i=1}^{n} \sum_{j:\hat{Y}_{i,j}(t)\geq 0} \Big(\Big| \mathrm{E}\left[\{\hat{Y}_{i,j}(t)\} - Z_{i,j}(t) \mid \mathbb{Y}_{\boldsymbol{\ell}}\right] \\ &\quad - \mathrm{E}\left[\{\hat{Y}_{i,j}(t)\} - Z_{i,j}(t) \mid \mathbb{Y}_{\boldsymbol{\ell-1}}\right] \Big| \Big) \cdot \left|\mathcal{C}^C_{k,i\to j}(t-s)\right| \\ &= \sum_{\{i,j\}\in E} \left| \mathrm{E}\left[Z^{(s)}_{i,j} \mid \mathbb{Y}_{\boldsymbol{\ell}}\right] - \mathrm{E}\left[Z^{(s)}_{i,j} \mid \mathbb{Y}_{\boldsymbol{\ell-1}}\right] \right| \\ &\quad \cdot \left|\mathcal{C}^C_{k,i\to j}(t-s)\right| \\ &= \sum_{\{i,j\}\in E} \left|\Lambda^{(s)}_{i,j}\right| \cdot \left|\mathcal{C}^C_{k,i\to j}(t-s)\right| \\ &\leq \sum_{i=1}^{n} \Big(\max_{j\in N(i)} \left|\mathcal{C}^C_{k,i\to j}(t-s)\right| \Big) \cdot \sum_{j\in N(i)} \left|\Lambda^{(s)}_{i,j}\right| \ , \end{aligned} \tag{8}$$

where we used

$$\Lambda^{(s)}_{i,j} = \mathrm{E}\left[Z^{(s)}_{i,j} \mid \mathbb{Y}_{\boldsymbol{\ell}}\right] - \mathrm{E}\left[Z^{(s)}_{i,j} \mid \mathbb{Y}_{\boldsymbol{\ell-1}}\right]$$

to simplify the notation.

As in [5], to bound (8) we consider $\sum_{\{i,j\}\in E} \left|\Lambda^{(s)}_{i,j}\right|$ for $i = i_1$ and $i \neq i_1$ separately.

Case 1: Let $i = i_1$. For each $j \in N(i_1)$, define indicator Bernoulli random variables $I_{u,j}, 1 \leq u \leq d$, where $I_{u,j}$ is one if the u'th excess token of i_1 in round s_1 goes to j and zero otherwise. Note that $Z^{(s_1)}_{i_1,j} = \sum_{1\leq u\leq d} I_{u,j}$. Let $r = \sum_{j\in N(i_1)} \left\{\hat{Y}_{i_1,j}(s_1)\right\}$ so that $\lceil r \rceil \geq b_1$ be the number of excess tokens of i_1 in round s_1. Clearly, r and the destinations of the excess tokens considered in the previous rounds, are

already determined by $Y_1, \ldots, Y_{\ell-1}$. The remaining receivers $Y_{\ell+1}, \ldots, Y_{\ell+r-b_1}$ are chosen independently from $N(i_1) \cup \{i_1\}$. Hence, the choice of Y_ℓ does not affect the distribution of $I_{u,j}$ except for $u = b_1$, and we have

$$\begin{aligned}
&\mathrm{E}\left[Z_{i_1,j}(s_1) \mid \mathbb{Y}_{\boldsymbol{\ell}}\right] - \mathrm{E}\left[Z_{i_1,j}(s_1) \mid \mathbb{Y}_{\boldsymbol{\ell}-1}\right] \\
&\quad = \mathrm{E}\left[I_{1,j} + \cdots + I_{\lceil r\rceil,j} \mid \mathbb{Y}_{\boldsymbol{\ell}}\right] \\
&\qquad - \mathrm{E}\left[I_{1,j} + \cdots + I_{\lceil r\rceil,j} \mid \mathbb{Y}_{\boldsymbol{\ell}-1}\right] \\
&\quad = \mathrm{E}\left[I_{b_1,j} \mid \mathbb{Y}_{\boldsymbol{\ell}}\right] - \mathrm{E}\left[I_{b_1,j} \mid \mathbb{Y}_{\boldsymbol{\ell}-1}\right]
\end{aligned}$$

Let $w \in N(i_1) \cup \{i_1\}$ be the destination of the b_1-th excess token of i_1 in round s_1, that is, $Y_\ell = w$ and hence,

$$\Lambda_{i_1,w}^{(s_1)} = 1 - \{\hat{Y}_{i_1,w}(t)\}/r \ .$$

For any $j \in N(i_1) \setminus \{w\}$ we have

$$\Lambda_{i_1,w}^{(s_1)} = -\{\hat{Y}_{i_1,j}(t)\}/r \ .$$

Hence,

$$\begin{aligned}
\sum_{j\in N(i_1)} \left|\Lambda_{i_1,j}^{(s_1)}\right| &\le 1 - \{\hat{Y}_{i_1,w}(t)\}/\lceil r\rceil \qquad (9)\\
&\quad + \sum_{j\in N(i_1)\setminus\{w\}} \{\hat{Y}_{i_1,j}(t)\}/\lceil r\rceil \\
&\le 1 + \sum_{j\in N(i_1)} \{\hat{Y}_{i_1,j}(t)\}/\lceil r\rceil \\
&\le 2, \qquad (10)
\end{aligned}$$

where the last inequality holds since

$$\sum_{j\in N(i_1)} \{\hat{Y}_{i_1,j}(t)\} = r \le \lceil r\rceil \ .$$

Case 2: $i \neq i_1$. As ℓ corresponds to (s_1, i_1, b_1), the random variable $Z_{i,j}(s_1)$ is independent of Y_ℓ when conditioned on $\mathbb{Y}_{\ell-1}$. Hence, similar to [5], we have

$$\begin{aligned}
\sum_{\{i,j\}\in E} \left|\Lambda_{i,j}^{(s_1)}\right| &= \sum_{j:\{i,j\}\in E} \left|\mathrm{E}\left[Z_{i,j}^{(s)} \mid \mathbb{Y}_{\boldsymbol{\ell}}\right] - \mathrm{E}\left[Z_{i,j}^{(s)} \mid \mathbb{Y}_{\boldsymbol{\ell}-1}\right]\right| \\
&= 0 \ .
\end{aligned}$$

Combining Case 1 and Case 2 we obtain

$$\begin{aligned}
&\left(\max_{j\in N(i_1)} \left|\mathcal{C}_{k,i\to j}^{C}(t-s)\right|\right) \sum_{j:\{i_1,j\}\in E} \left|\Lambda_{i_1,j}^{(s)}\right| \\
&+ \sum_{i\in V, i\neq i_1} \left(\max_{j\in N(i)} \left|\mathcal{C}_{k,i\to j}^{C}(t-s)\right|\right) \sum_{\{i,j\}\in E} \left|\Lambda_{i,j}^{(s)}\right| \\
&\le \max_{j\in N(i_1)} \left|\mathcal{C}_{k,i\to j}^{C}(t-s)\right| \cdot 2 + 0 \ . \qquad (11)
\end{aligned}$$

Case $\mathbf{s > s_1}$: Let $\widetilde{\ell}$ be the largest integer that corresponds to round $s-1$. Since $s > s_1$, we have $s - 1 \ge s_1$ and therefore $\widetilde{\ell} \ge \ell$. By the choice of $\widetilde{\ell}$, $Y_{\widetilde{\ell}}, \ldots, Y_1$ determine the load vector at the end of round s_1, $X^{(s_1)}$. By Observation 1, we obtain $\mathrm{E}\left[E_{i,j}(s) \mid Y_{\widetilde{\ell}}, \ldots, Y_1\right] = 0$, and by the law of total expectation,

$$\begin{aligned}
&\mathrm{E}\left[E_{i,j}(s) \mid \mathbb{Y}_{\boldsymbol{\ell}}\right] \\
&\quad = \mathrm{E}\left[\mathrm{E}\left[E_{i,j}(s) \mid Y_{\widetilde{\ell}}, \ldots, Y_1\right] \mid Y_\ell, Y_{\ell-1}, \ldots, Y_1\right] \\
&\quad = \mathrm{E}\left[0 \mid Y_\ell, Y_{\ell-1}, \ldots, Y_1\right] = 0 \ .
\end{aligned}$$

With the same arguments, $\mathrm{E}\left[E_{i,j}(s) \mid \mathbb{Y}_{\boldsymbol{\ell}-1}\right] = 0$, and thus

$$\begin{aligned}
&\sum_{s=s_1+1}^{t} \sum_{\{i,j\}\in E} \left|\mathrm{E}\left[E_{i,j}(s) \mid \mathbb{Y}_{\boldsymbol{\ell}}\right] - \mathrm{E}\left[E_{i,j}(s) \mid \mathbb{Y}_{\boldsymbol{\ell}-1}\right]\right| \\
&\cdot \left|\mathcal{C}_{k,i\to j}^{C}(t-s)\right| = 0 \ . \qquad (12)
\end{aligned}$$

This finishes the case distinction. Combining equations (7), (11), and (12) for the three cases $s < s_1$, $s = s_1$, and $s > s_1$, similar to [5] we obtain that for every fixed $1 \le \ell \le tnd$,

$$\begin{aligned}
c_\ell &= \left|\mathrm{E}\left[f_k \mid \mathbb{Y}_{\boldsymbol{\ell}}\right] - \mathrm{E}\left[f_k \mid \mathbb{Y}_{\boldsymbol{\ell}-1}\right]\right| \\
&\le \sum_{s=0}^{t} \sum_{\{i,j\}\in E} \left|\mathrm{E}\left[E_{i,j}(s) \mid \mathbb{Y}_{\boldsymbol{\ell}}\right] - \mathrm{E}\left[E_{i,j}(s) \mid \mathbb{Y}_{\boldsymbol{\ell}-1}\right]\right| \\
&\quad \cdot \left|\mathcal{C}_{k,i_1\to j}^{C}(t-s_1)\right| \\
&= 0 + \max_{j\in N(i_1)} \left|\mathcal{C}_{k,i_1\to j}^{C}(t-s_1)\right| \cdot 2 + 0 \\
&= 2 \cdot \max_{j\in N(i_1)} \left|\mathcal{C}_{k,i_1\to j}^{C}(t-s_1)\right|.
\end{aligned}$$

Now we consider the sum of the error terms.

$$\begin{aligned}
\sum_{\ell=1}^{(t+1)nd} (c_\ell)^2 &\le \sum_{s=0}^{t} \sum_{i=1}^{n} \sum_{b=1}^{d} \left(2 \max_{j\in N(i)} \left|\mathcal{C}_{k,i\to j}^{C}(t-s)\right|\right)^2 \\
&= 4d \sum_{s=0}^{t} \sum_{i=1}^{n} \max_{j\in N(i)} \left(\mathcal{C}_{k,i\to j}^{C}(s)\right)^2 \\
&\le 4d \max_{k\in V} \left(\sum_{s=0}^{\infty} \sum_{i=1}^{n} \max_{j\in N(i)} \left(\mathcal{C}_{k,i\to j}^{C}(s)\right)^2\right) \\
&= 8d \left(\Upsilon^{\mathrm{C}}(G)\right)^2 \ . \qquad (13)
\end{aligned}$$

So by Azuma's inequality [9, p. 68] we have for any $\delta \ge 0$,

$$\mathrm{P_r}\left[|f_k| > \delta\right] \le 2 \exp\left(-\delta^2 / \left(2 \sum_{\ell=1}^{tnd} (c_\ell)^2\right)\right) \ .$$

Hence by choosing $\delta = \Upsilon^C(G)\sqrt{32d \ln n}$, the probability above gets smaller than $2\,n^{-2}$. Applying the union bound we obtain

$$\mathrm{P_r}\left[\exists k \in V \colon |f_k| > \delta\right] \le 2n^{-1} \ .$$

This implies

$$\mathrm{P_r}\left[\max_{i,j\in[n]} \left|X_i^R(t) - x_i^C(t)\right| \le \delta\right] \ge 1 - 2n^{-1} \ ,$$

which finishes the proof. ∎

Using Theorem 3 we can also obtain concrete results for randomized FOS processes as stated in the following theorems. The first result holds for the homogeneous case and a special class of algorithms where $\alpha_{i,j} = 1/(\gamma d)$ only. Recall that d is the maximum degree. The same result was already shown in [21].

Observation 3. *Assume $s_1 = s_2 = \ldots = s_n$ and $\alpha_{i,j} = \frac{1}{\gamma d}$. Let C be a continuous FOS process and let $R = R(C)$ be a discrete FOS process based on the rounding algorithm applied on C. Then*

(1) $\Upsilon^{\mathrm{C}}(G) = O\left(\sqrt{\gamma d/(2-2/\gamma)}\right)$.

(2) For any round t we have w.h.p.

$$\left|x_k^R(t) - x_k^{\mathrm{C}}(t)\right| = O\left(\sqrt{\frac{\gamma d}{2-2/\gamma}} \cdot \sqrt{d \log n}\right) \ .$$

In [21] the authors applied a potential function in order to estimate $\Upsilon^C(G)$. This proof relies heavily on the fact that the transition probabilities are uniform for all edges, which is not the case for the heterogeneous model or the case where the $\alpha_{i,j}$ depend on d_i and d_j only. The next result is more general and applies to both of these cases as well.

Theorem 4. *Let C be a continuous FOS process and let $R = R(C)$ be a discrete FOS process based on the rounding algorithm applied on C. Then*

(1) $\Upsilon^{\mathrm{C}}(G) = O\left(\sqrt{d \cdot \log s_{\max}/(1-\lambda)}\right)$.

(2) For any round t we have w.h.p.

$$\left|x_k^R(t) - x_k^{\mathrm{C}}(t)\right| = O\left(d \cdot \sqrt{\frac{\log n \cdot \log s_{\max}}{1-\lambda}}\right) \ .$$

To show Theorem 4 we first show the following lemma.

Lemma 5. *For an arbitrary $1 \le k \le n$, let the vector $\mathbf{a}$ be such that $\mathbf{a}_i = M^t_{k,i} - \frac{s_k}{s}$. Then we have*

$$\|\mathbf{a}\|_2^2 \le 2\, s_{\max}\, \lambda^{2t}$$

Proof: Let $\hat{\mathbf{a}} = \hat{\mathbf{k}} - \frac{s_k}{s} \cdot \mathbf{1}_n$. Note that $\mathbf{a} = \hat{\mathbf{a}}\, M^t$. Let $\mathbf{v}_1, \ldots, \mathbf{v}_n$ be the eigenvectors of M^t with eigenvalues $\lambda_1^t, \ldots, \lambda_n^t$, and λ^t be the second largest eigenvalue. Using the fact that $M = I - LS^{-1}$, it is not hard to see that $S^{-1}M^t$ is symmetric. Hence, for each right eigenvector $\mathbf{v}_i$ of M^t there is a left eigenvector $\mathbf{u}_i = S^{-1}\mathbf{v}_i$ with the same eigenvalue λ_i^t as proved in the following.

$$\begin{aligned}(M^t)^T\mathbf{u}_i &= (M^t)^T S^{-1} S\mathbf{u}_i = (S^{-1}M^t)^T\mathbf{v}_i \\ &= S^{-1}M^t\mathbf{v}_i = \lambda_i^t S^{-1}\mathbf{v}_i = \lambda_i^t\mathbf{u}_i\end{aligned}$$

Also, note that $S^{-1}M^tS\mathbf{u}_i = S^{-1}M^t\mathbf{v}_i = \lambda_i^t S^{-1}\mathbf{v}_i = \lambda_i^t\mathbf{u}_i$. As a result, $\mathbf{u}_i$'s are eigenvectors of $S^{-1}M^tS$, which is symmetric because it is the product of symmetric matrices $S^{-1}M^t$ and S. Therefore, $\mathbf{u}_1, \ldots, \mathbf{u}_n$ form an orthonormal basis; so we can write $\hat{\mathbf{a}} = \sum_{i=1}^n c_i\mathbf{u}_i$. Now we write

$$\mathbf{a} = \hat{\mathbf{a}}\, M^t = \sum_{i=1}^n c_i\mathbf{u}_i M^t = \sum_{i=1}^n \lambda_i^t c_i \mathbf{u}_i$$

Therefore,

$$\|\mathbf{a}\|_2^2 = \sum_{i=1}^n \lambda_i^{2t} c_i^2 \|\mathbf{u}_i\|_2^2 \le \lambda^{2t}\sum_{i=1}^n c_i^2\|\mathbf{u}_i\|_2^2 = \lambda^{2t}\|\hat{\mathbf{a}}\|_2^2 \quad (14)$$

where the inequality uses the fact that $\hat{\mathbf{a}} \perp (s_1, \ldots, s_n)$ which is the eigenvector corresponding to the largest eigenvalue. Also, the last equality follows from the fact that $\mathbf{u}_i$'s form an orthonormal basis. On the other hand,

$$\begin{aligned}\|\hat{\mathbf{a}}\|_2^2 &\le n \cdot \frac{s_k^2}{s^2} + 1 \le \frac{n\, s_k^2}{(n-1+s_k)^2} + 1 \\ &\le \frac{n\, s_k^2}{2(n-1)s_k} + 1 \le s_k + 1 \le 2 s_{\max}\end{aligned}$$

Together with (14), this yields

$$\|\mathbf{a}\|_2^2 \le 2\, s_{\max}\lambda^{2t} \ ,$$

as required. ∎

Proof of Theorem 4: We have

$$\begin{aligned}\left(\Upsilon^{\mathrm{FOS}}(G)\right)^2 &= \sum_{t=0}^{\infty}\sum_{i=1}^{n}\max_{j\in N(i)}\left(M^t_{k,i} - M^t_{k,j}\right)^2 \\ &\le \sum_{t=0}^{\infty}\sum_{i=1}^{n}\sum_{j\in N(i)}\left(M^t_{k,i} - M^t_{k,j}\right)^2 \\ &= \sum_{t=0}^{t_1-1}\sum_{i=1}^{n}\sum_{j\in N(i)}\left(M^t_{k,i} - M^t_{k,j}\right)^2 \\ &\quad + \sum_{t=t_1}^{\infty}\sum_{i=1}^{n}\sum_{j\in N(i)}\left(M^t_{k,i} - M^t_{k,j}\right)^2 \quad (15)\end{aligned}$$

In the following, we use $\sigma = \beta - 1$ for brevity. Note that $0 < \sigma < 1$.

$$\begin{aligned}&\sum_{t=0}^{t_1-1}\sum_{i=1}^{n}\sum_{j\in N(i)}\left(M^t_{k,i} - M^t_{k,j}\right)^2 \\ &\le \sum_{t=0}^{t_1}\sum_{i=1}^{n}\sum_{j\in N(i)} 2\left((M^t_{k,i})^2 + (M^t_{k,j})^2\right) \\ &\le 4\cdot d\cdot\sum_{t=0}^{t_1-1}\sum_{i=1}^{n}(M^t_{k,i})^2 \\ &\le 4\cdot d\cdot\sum_{t=0}^{t_1-1}\|M^t\,\hat{\mathbf{k}}\|_2^2 \\ &= 4\cdot d\cdot\sum_{t=0}^{t_1-1}\left(\|\hat{\mathbf{k}}\|_2\cdot\max_i\lambda_i^t\right)^2 \\ &= 4\cdot d\cdot\sum_{t=0}^{t_1-1}1 \le 4\cdot d\cdot t_1 \quad (16)\end{aligned}$$

Let $t_1 = (\log s_{\max})/(2-2\lambda)$. Note that $\lambda^{1/(1-\lambda)} \le 1/e$. Then we have

$$\begin{aligned}&\sum_{t=t_1}^{\infty}\sum_{i=1}^{n}\sum_{j\in N(i)}\left(M^t_{k,i} - M^t_{k,j}\right)^2 \\ &\le \sum_{t=t_1}^{\infty}\sum_{i=1}^{n}\sum_{j\in N(i)} 2\left(\left(M^t_{k,i} - \frac{s_k}{s}\right)^2 + \left(M^t_{k,j} - \frac{s_k}{s}\right)^2\right) \\ &= 4\cdot d\cdot\sum_{t=t_1}^{\infty}\sum_{i=1}^{n}\left(M^t_{k,i} - \frac{s_k}{s}\right)^2 \\ &\le 8\cdot d\cdot s_{\max}\cdot\sum_{t=t_1}^{\infty}\lambda^{2t} \quad (17) \\ &\le 8\cdot d\cdot s_{\max}\cdot\lambda^{2t_1}\cdot\frac{1}{1-\lambda} \le \frac{8\,d}{1-\lambda} \quad (18)\end{aligned}$$

where (17) follows from Lemma 5. Combining equations (15), (18), and (16) we get

$$\left(\Upsilon^{\mathrm{FOS}}(G)\right)^2 = O\left(\frac{d\cdot\log s_{\max}}{1-\lambda}\right) \ , \quad (19)$$

which proves the first statement. The bound in the second statement follows immediately from statement (1) and Theorem 3. ■

IV. Second Order Diffusion Processes

In this section we show that after some slight adjustments the framework of Section III can be applied to second order processes on heterogeneous networks. All we have to do is to state definitions 2 and 3 in a more general way that captures the dependence of SOS on the load transfer of the previous round. It is easy to see that Lemma 2 and Theorem 3 still hold assuming the new definitions. (Note that SOS is linear). If C is a second order process, then $y^C(t)$ is determined based on $x^C(t)$ and $y^C(t-1)$. More formally, $y^C(t) = C(x^C(t), y^C(t-1))$. Thus, the new definitions also incorporate $y^C(t-1)$. We again use $\hat{\mathbf{i}}$ to denote the unit vector with 1 as its i'th entry.

Definition 4 (Linearity)**.** *A process A is said to be* linear *if for all $\mathbf{x}, \mathbf{x}' \in \mathbb{R}^n, \mathbf{y}, \mathbf{y}' \in \mathbb{R}^{n\times n}$ and $a, b \in \mathbb{R}$ we have $A(a\mathbf{x} + b\mathbf{x}', a\mathbf{y} + b\mathbf{y}') = aA(\mathbf{x}, \mathbf{y}) + bA(\mathbf{x}', \mathbf{y}')$.*

Definition 5 (Contributions)**.** *Let $\mathbf{x}(0) = \mathbf{x}'(0) = \hat{\mathbf{i}}$, $\mathbf{y}(0) = \mathbf{0}_{n\times n}$ and let $\mathbf{y}'(0)$ be also all zero except $\mathbf{y}'_{i,j}(0) = 1$, so that $\mathbf{x}(1) = \hat{\mathbf{i}}$, $\mathbf{x}'(1) = \hat{\mathbf{j}}$. Let $\mathbf{x}(t+1)$ and $\mathbf{x}'(t+1)$ be the load vectors obtained from applying C for t rounds on $(\mathbf{x}(1), \mathbf{y}(0))$ and $(\mathbf{x}'(1), \mathbf{y}'(0))$, respectively. Then the* contribution *of the edge (i, j) on a node k after t rounds is defined as $\mathcal{C}^C_{k,i\to j}(t) = \mathbf{x}_k(t) - \mathbf{x}'_k(t)$.*

To prove bounds of the deviation of theorems 8 and 9 we apply Observation 4 which follows from Lemma 2 and Theorem 3. This gives us an upper bound in terms of the $\mathcal{C}^C_{k,i\to j}(t)$'s. Hence to obtain a more concrete bound we have to estimate $\mathcal{C}^C_{k,i\to j}(t)$ which is done in Lemma 6 and upper bounded in Lemma 7. The contributions are expressed based on a sequence of matrices $Q(t)$ defined below, whose role in error propagation is similar to that of the diffusion matrix in FOS.

$$Q(t) = \begin{cases} \mathbf{I} & \text{if } t = 0 \\ \beta \cdot M & \text{if } t = 1 \\ \beta \cdot M\, Q(t-1) + (1-\beta) \cdot Q(t-2) & \text{if } t \geq 2 \end{cases} \tag{20}$$

Lemma 6. *For $t > 0$, we have*

$$\mathcal{C}^{\mathrm{SOS}}_{k,i\to j}(t) = Q_{k,i}(t-1) - Q_{k,j}(t-1)\ .$$

Proof: Let $\mathbf{x}(0) = \mathbf{x}'(0) = \hat{\mathbf{i}}$, $\mathbf{y}(0) = \mathbf{0}_{n\times n}$ and let $\mathbf{y}'(0)$ be also all zero except $\mathbf{y}'_{i,j}(0) = 1$, so that $\mathbf{x}(1) = \hat{\mathbf{i}}$, $\mathbf{x}'(1) = \hat{\mathbf{j}}$. Let $\mathbf{x}(t+1)$ and $\mathbf{x}'(t+1)$ be the load vectors obtained from applying SOS for t rounds on $(\mathbf{x}(1), \mathbf{y}(0))$ and $(\mathbf{x}'(1), \mathbf{y}'(0))$, respectively. Let $\mathcal{C}^{\mathrm{SOS}}_{i\to j}(t)$ be a vector that has $\mathcal{C}^{\mathrm{SOS}}_{k,i\to j}(t)$ as its k'th entry, for $1 \leq k \leq n$. Let $\mathbf{w} = \hat{\mathbf{i}} - \hat{\mathbf{j}}$. Then we have

$$\mathcal{C}^{\mathrm{SOS}}_{i\to j}(t) = \mathbf{x}(t) - \mathbf{x}'(t) = \begin{cases} \mathbf{0} & \text{if } t = 0 \\ \mathbf{w} & \text{if } t = 1 \\ \beta \cdot M\, \mathcal{C}^{\mathrm{SOS}}_{i\to j}(t-1) & \\ \quad +(1-\beta) \cdot \mathcal{C}^{\mathrm{SOS}}_{i\to j}(t-2) & \text{if } t \geq 2 \end{cases}$$

where the third equation holds because for all $t \geq 2$, both $\mathbf{x}(t)$ and $\mathbf{x}'(t)$ follow the same equation $x(t) = \beta \cdot Mx(t-1) + (1-\beta) \cdot x(t-2)$. Now, it can be proved by induction that $\mathcal{C}^{\mathrm{SOS}}_{i\to j}(t) = Q(t-1)\,\mathbf{w}$. Recall that all entries of $\mathbf{w}$ are zero except $\mathbf{w}_i = 1$ and $\mathbf{w}_j = -1$. Therefore, for $t > 0$ we get $\mathcal{C}^{\mathrm{SOS}}_{k,i\to j}(t) = Q_{k,i}(t-1) - Q_{k,j}(t-1)$. ■

The following lemma provides a bound for the second norm of $Q(t)$, which is later used in the proofs of theorems 8 and 9.

Lemma 7. *Let $\beta = \beta_{opt} = 2/(1 + \sqrt{1-\lambda^2})$. The following statements are true.*

1) *Eigenvectors of $Q(t)$ form a basis for $\mathbb{R}^n$.*
2) *Let $\gamma = \left(\sqrt{\beta - 1}\right)^t (t+1)$. Then γ is an upper bound on the eigenvalues of $Q(t)$ except the eigenvalue corresponding to the eigenvector $(s_1, \cdots, s_n)$.*
3) *$Q(t)$ has equal column sums.*
4) *Define $q(t) = \sum_{1\leq j\leq n} Q_{i,j}(t)$ for an arbitrary $1 \leq i \leq n$ (note that by the statement (3), this is a valid definition). Fix a $1 \leq k \leq n$, and let the vector $\mathbf{a}$ be such that $\mathbf{a}_i = Q_{k,i}(t) - s_k/s \cdot q(t)$. Then we have $\|\mathbf{a}\|_2^2 \leq 2\, s_{\max}(\beta-1)^t(t+1)^2$.*

Proof:

a) Proof of (1): First we observe that the eigenvectors of $Q(t)$ are the same as the eigenvectors of M. This can be proved by an induction using the recurrence of (20) as follows.

Suppose $\mathbf{v}$ is an eigenvector of M with eigenvalue α. Then $\mathbf{v}$ is also an eigenvector of $Q(t)$ and $Q(t-1)$ by the induction hypothesis. Let μ_1 and μ_2 be the corresponding eigenvalues. We have

$$\begin{aligned} Q(t+1)\mathbf{v} &= \beta \cdot M\, Q(t)\mathbf{v} + (1-\beta) \cdot Q(t-1)\mathbf{v} \\ &= \beta\mu_1 \cdot M\,\mathbf{v} + (1-\beta)\mu_2\,\mathbf{v} \\ &= \beta\mu_1\alpha\mathbf{v} + (1-\beta)\mu_2\mathbf{v} \\ &= (\beta\mu_1\alpha + (1-\beta)\mu_2)\,\mathbf{v}, \end{aligned}$$

which shows that $\mathbf{v}$ is also an eigenvector of $Q(t+1)$. Also, note that $M = I - LS^{-1}$ where L is the Laplacian matrix of the graph and S is the diagonal matrix of speeds. The eigenvectors of M are the same as those of LS^{-1}. By [11, proof of Lemma 1] the eigenvectors of LS^{-1} form a basis for $\mathbb{R}^n$. Therefore the eigenvectors of M and the eigenvectors of $Q(t)$ form a basis for $\mathbb{R}^n$.

b) Proof of (2): From the induction in the proof of statement (1) one can see that corresponding to each eigenvalue λ_j of M an eigenvalue $\gamma_j(t)$ of $Q(t)$ can be obtained according to the following recursion.

$$\gamma_j(t) = \begin{cases} 1 & \text{if } t = 0 \\ \beta\lambda_j & \text{if } t = 1 \\ \beta\lambda_j \cdot \gamma_j(t-1) + (1-\beta) \cdot \gamma_j(t-2) & \text{if } t \geq 2 \end{cases}$$

Solving the above recursion we get

$$\gamma_j(t) = \begin{cases} \frac{1-(\beta-1)^{t+1}}{2-\beta} & \text{if } \lambda_j = 1 \\ \left(\sqrt{\beta-1}\right)^t (t+1) & \text{if } |\lambda_j| = \lambda \\ r^t \left(\cos(\theta t) + \sin(\theta t) \cdot \frac{\lambda_j}{\sqrt{\lambda^2 - \lambda_j^2}}\right) & \text{if } |\lambda_j| < \lambda \end{cases}$$

where $r = \sqrt{\beta - 1}$, and $0 < \theta < \pi$ is such that $\sin\theta = \sqrt{\lambda^2 - \lambda_j^2}/\lambda$, and $\cos\theta = \lambda_j/\lambda$. Note that the eigenvalue corresponding to $\lambda_j = 1$ belongs to the eigenvector $(s_1, \cdots, s_n)$. Hence, it suffices to prove that in (IV-0b) the case $|\lambda_j| < \lambda$ does not produce eigenvalues bigger than those obtained in the case $|\lambda_j| = \lambda$. Note that

$$\begin{aligned}\gamma_j(t) &= r^t \left(\cos(\theta t) + \sin(\theta t) \cdot \frac{\lambda_j}{\sqrt{\lambda^2 - \lambda_j^2}}\right) \\ &\le \left(\sqrt{\beta-1}\right)^t \cdot \frac{\sin((t+1)\theta)}{\sin\theta} \qquad (21) \\ &\le \left(\sqrt{\beta-1}\right)^t \cdot (t+1),\end{aligned}$$

where in (21) we use $\sin(nx) \le n \sin x$ for $0 < x < \pi$ and $n \in \mathbb{N}$.

c) Proof of (3): The statement follows from a simple induction using (20). The case $Q(0) = \mathbf{I}$ is trivial. $Q(1) = \beta M$ also has equal column sums, since the entries in each column of M sum to one (this is necessay to guarantee load conservation). Suppose for all $t_1 \le t$, $Q(t_1)$ has equal column sums. Let us denote this value by $q(t_1)$.

$$\begin{aligned}\mathbf{1}_n Q(t+1) &= \beta \cdot \mathbf{1}_n M\, Q(t) + (1-\beta) \cdot \mathbf{1}_n Q(t-1) \\ &= \beta \cdot \mathbf{1}_n Q(t) + (1-\beta) \cdot q(t-1) \cdot \mathbf{1}_n \\ &= \beta \cdot q(t) \cdot \mathbf{1}_n + (1-\beta) \cdot q(t-1) \cdot \mathbf{1}_n \\ &= (\beta \cdot q(t) + (1-\beta) \cdot q(t-1)) \cdot \mathbf{1}_n\end{aligned}$$

which shows that all column sums of $Q(t+1)$ are equal to $\beta \cdot q(t) + (1-\beta) \cdot q(t-1)$.

d) Proof of (4): Let $\hat{\mathbf{a}} = \hat{\mathbf{k}} - \frac{s_k}{s} \cdot \mathbf{1}_n$. Note that $\mathbf{a} = \hat{\mathbf{a}}\, Q(t)$. Let $\mathbf{v}_1, \ldots, \mathbf{v}_n$ be the eigenvectors of $Q(t)$ with eigenvalues $\gamma_1, \ldots, \gamma_n$, and γ be defined as in the statement (2) of the lemma. Using the fact that $M = I - LS^{-1}$, it can be proved by induction that $S^{-1}Q(t)$ is symmetric. Hence, for each right eigenvector $\mathbf{v}_i$ of $Q(t)$ there is a left eigenvector $\mathbf{u}_i = S^{-1}\mathbf{v}_i$ with the same eigenvalue γ_i as proved in the following.

$$\begin{aligned}(Q(t))^T \mathbf{u}_i &= (Q(t))^T S^{-1} S \mathbf{u}_i = (S^{-1}Q(t))^T \mathbf{v}_i \\ &= S^{-1}Q(t)\mathbf{v}_i = \gamma_i S^{-1}\mathbf{v}_i = \gamma_i \mathbf{u}_i\end{aligned}$$

Also, note that $S^{-1}Q(t)S\mathbf{u}_i = S^{-1}Q(t)\mathbf{v}_i = \gamma_i S^{-1}\mathbf{v}_i = \gamma_i \mathbf{u}_i$. As a result, $\mathbf{u}_i$'s are eigenvectors of $S^{-1}Q(t)S$, which is symmetric because it is the product of symmetric matrices $S^{-1}Q(t)$ and S. Therefore, $\mathbf{u}_1, \ldots, \mathbf{u}_n$ form an orthonormal basis; so we can write $\hat{\mathbf{a}} = \sum_{i=1}^n c_i \mathbf{u}_i$. Now we write

$$\mathbf{a} = \hat{\mathbf{a}}\, Q(t) = \sum_{i=1}^n c_i \mathbf{u}_i Q(t) = \sum_{i=1}^n \gamma_i c_i \mathbf{u}_i \ .$$

Therefore,

$$\|\mathbf{a}\|_2^2 = \sum_{i=1}^n \gamma_i^2 c_i^2 \|\mathbf{u}_i\|_2^2 \le \gamma^2 \sum_{i=1}^n c_i^2 \|\mathbf{u}_i\|_2^2 = \gamma^2 \|\hat{\mathbf{a}}\|_2^2 \qquad (22)$$

where the inequality uses the fact that $\hat{\mathbf{a}} \perp (s_1, \ldots, s_n)$ and part (2) of the lemma, and the last equality follows from the fact that $\mathbf{u}_i$'s form an orthonormal basis. Also,

$$\begin{aligned}\|\hat{\mathbf{a}}\|_2^2 &\le n \cdot \frac{s_k^2}{s^2} + 1 \\ &\le \frac{n\, s_k^2}{(n - 1 + s_k)^2} + 1 \\ &\le \frac{n\, s_k^2}{2(n-1)s_k} + 1 \\ &\le s_k + 1 \\ &\le 2 s_{\max} \ .\end{aligned}$$

Together with (22), this yields $\|\mathbf{a}\|_2^2 \le 2\, s_{\max} (\beta-1)^t (t+1)^2$. ■

A. Deviation between Continuous and Discrete SOS

In this section we show a bound on the deviation between a continuous SOS and its rounded version. The authors of [12] show a similar bound on the deviation using the second norm, i.e., they show a bound of $||x^{D(\mathrm{SOS})}(t) - x^{\mathrm{SOS}}(t)||_2 = O\left(d\sqrt{n s_{\max}}/(1-\lambda)\right)$. Note that the bound on the deviation for FOS, which is $O(d\sqrt{s_{\max} \log n/(1-\lambda)})$, is smaller.

Theorem 8. *Consider a discrete SOS process $D = D(\mathrm{SOS})$ with optimal β and a rounding scheme that rounds a fractional value to either its floor or its ceiling. Then for arbitrary $t \ge 0$ we have $\left|x_k^D(t) - x_k^{\mathrm{SOS}}(t)\right| = O\left(d\sqrt{n s_{\max}}/(1-\lambda)\right)$.*

Proof: We use Lemma 2 to obtain a deviation bound for the general case. We have

$$\begin{aligned}&\left|x_k^{D(\mathrm{SOS})}(t+1) - x_k^{\mathrm{SOS}}(t+1)\right| \\ &\quad= \left|\sum_{s=0}^{t} \sum_{\{i,j\} \in E} (Q_{k,i}(s) - Q_{k,j}(s)) \cdot e_{i,j}(t-s)\right| \\ &\quad\le \sum_{s=0}^{t} \sum_{\{i,j\} \in E} |Q_{k,i}(s) - Q_{k,j}(s)| \\ &\quad\le \sum_{s=0}^{t} \sum_{\{i,j\} \in E} \left(\left|Q_{k,i}(s) - \frac{s_k}{s} \cdot q(s)\right| \right. \\ &\qquad\left. + \left|Q_{k,j}(s) - \frac{s_k}{s} \cdot q(s)\right|\right) \\ &\quad= d \cdot \sum_{s=0}^{t} \sum_{i=1}^{n} \left|Q_{k,i}(s) - \frac{s_k}{s} \cdot q(s)\right| \\ &\quad\le d \cdot \sqrt{n} \cdot \sum_{s=0}^{t} \left(\sum_{i=1}^{n} \left(Q_{k,i}(s) - \frac{s_k}{s} \cdot q(s)\right)^2\right)^{1/2} \qquad (23) \\ &\quad\le 4 \cdot d \cdot \sqrt{n} \cdot \sqrt{2 s_{\max}} \cdot \sum_{s=0}^{\infty} \left(\sqrt{\beta-1}\right)^s (s+1) \qquad (24) \\ &\quad\le 4 \cdot d \cdot \sqrt{2 n s_{\max}} \cdot \frac{1}{\left(1 - \sqrt{\beta-1}\right)^2} \\ &\quad\le 16 \cdot d \cdot \sqrt{2 n s_{\max}} \cdot \frac{1}{1-\lambda}\end{aligned}$$

where Equation 23 follows from the Cauchy-Schwarz inequality and Equation 24 follows from Lemma 7.(4).

■

B. Framework for Randomized SOS

In the next theorem we bound the deviation between continuous and discrete SOS using the randomized rounding scheme from Section III-B. As mentioned earlier in this section, is easy to see that the proof of Theorem 3 holds for the more general definitions of linearity and contribution of this section. Hence, we can state the following observation and show similar to Section III the next theorem.

Observation 4. *In the setting of Section IV for an arbitrary round t we have w.h.p.*

$$\left|X_k^R(t) - x_k^C(t)\right| = O\left(\Upsilon^C(G) \cdot \sqrt{d \log n}\right)$$

Similar to Section III, we can use Observation 4 to show the next theorem.

Theorem 9. *Let $R = R(\mathrm{SOS})$ be a randomized-rounding discrete SOS process with optimal β obtained using our randomized rounding scheme. Then*

(1) $\Upsilon^{\mathrm{SOS}}(G) = O\left(\sqrt{d} \cdot \log s_{\max}/(1-\lambda)^{3/4}\right)$.

(2) The deviation of R from the continuous SOS in an arbitrary round t is w.h.p.

$$\left|x_k^R(t) - x_k^{\mathrm{SOS}}(t)\right| = O\Big(\frac{d \cdot \log s_{\max} \cdot \sqrt{\log n}}{(1-\lambda)^{3/4}}\Big) \ .$$

Proof: The bound on the refined local divergence is obtained using the formulation of Lemma 6 and the bound in Lemma 7. This bound together with the parametric deviation bound of Theorem 3 yield the second statement of the theorem.

We write

$$\begin{aligned}
\left(\Upsilon^{\mathrm{SOS}}(G)\right)^2 &= \sum_{t=0}^{\infty}\sum_{i=1}^{n} \max_{j\in N(i)} \left(Q_{k,i}(t) - Q_{k,j}(t)\right)^2 \\
&\le \sum_{t=0}^{\infty}\sum_{i=1}^{n} \sum_{j\in N(i)} \left(Q_{k,i}(t) - Q_{k,j}(t)\right)^2 \\
&= \sum_{t=0}^{t_1-1}\sum_{i=1}^{n} \sum_{j\in N(i)} \left(Q_{k,i}(t) - Q_{k,j}(t)\right)^2 \\
&\quad + \sum_{t=t_1}^{\infty}\sum_{i=1}^{n} \sum_{j\in N(i)} \left(Q_{k,i}(t) - Q_{k,j}(t)\right)^2
\end{aligned}$$

In the following, we use $\sigma = \beta - 1$ for brevity. Note that we have $0 < \sigma < 1$.

$$\begin{aligned}
&\sum_{t=0}^{t_1-1}\sum_{i=1}^{n} \sum_{j\in N(i)} \left(Q_{k,i}(t) - Q_{k,j}(t)\right)^2 \\
&\quad \le \sum_{t=0}^{t_1}\sum_{i=1}^{n} \sum_{j\in N(i)} 2\left((Q_{k,i}(t))^2 + (Q_{k,j}(t))^2\right) \\
&\quad \le 4\cdot d \cdot \sum_{t=0}^{t_1-1}\sum_{i=1}^{n} (Q_{k,i}(t))^2 \le 4\cdot d \cdot \sum_{t=0}^{t_1-1} \|Q(t)\,\hat{\mathbf{k}}\|_2^2 \\
&\quad = 4\cdot d \cdot \sum_{t=0}^{t_1-1} \left(\|\hat{\mathbf{k}}\|_2 \cdot \max_i \gamma_i(t)\right)^2 \\
&\quad = 4\cdot d \cdot \sum_{t=0}^{t_1-1} \left(\frac{1-\sigma^{t+1}}{1-\sigma}\right)^2 \\
&\quad \le 4\cdot d \cdot t_1 \cdot \left(\frac{1-\sigma^{t_1}}{1-\sigma}\right)^2 \le 4\cdot d \cdot t_1 \cdot (1-\sigma)^{-2} \qquad (25)
\end{aligned}$$

We also get

$$\begin{aligned}
&\sum_{t=t_1}^{\infty}\sum_{i=1}^{n} \sum_{j\in N(i)} \left(Q_{k,i}(t) - Q_{k,j}(t)\right)^2 \\
&\quad \le \sum_{t=t_1}^{\infty}\sum_{i=1}^{n} \sum_{j\in N(i)} 2\left(Q_{k,i}(t) - \frac{s_k}{s}\cdot q(t)\right)^2 \\
&\qquad + 2\left(Q_{k,j}(t) - \frac{s_k}{s}\cdot q(t)\right)^2 \\
&\quad = 4\cdot d \cdot \sum_{t=t_1}^{\infty}\sum_{i=1}^{n} \left(Q_{k,i}(t) - \frac{s_k}{s}\cdot q(t)\right)^2 \qquad (26) \\
&\quad \le 8\cdot d \cdot s_{\max} \cdot \sum_{t=t_1}^{\infty} \left(\sigma^t\,(t+1)^2\right) \qquad (27)
\end{aligned}$$

where the last inequality follows from part (4) of Lemma 7. The above summation can be bounded as follows.

$$\begin{aligned}
\sum_{t=t_1}^{\infty} \left(\sigma^t\,(t+1)^2\right) &= \frac{d}{d\sigma}\left(\sigma \frac{d}{d\sigma}\left(\frac{\sigma^{t_1+1}}{1-\sigma}\right)\right) \\
&\le \frac{2\sigma^{t_1+2}}{(1-\sigma)^3} + \frac{(2t_1+3)\cdot\sigma^{t_1+1}}{(1-\sigma)^2} + \frac{(t_1+1)^2\cdot\sigma^{t_1}}{1-\sigma} \qquad (28)
\end{aligned}$$

Let $t_1 = (\log s_{\max})/(1-\sigma)$. Note that $\sigma^{1/(1-\sigma)} \le 1/e$. Then (28) yields

$$\sum_{t=t_1}^{\infty} \left(\sigma^t\,(t+1)^2\right) = O\left(\frac{\log^2 s_{\max}}{s_{\max}\cdot(1-\sigma)^3}\right) \ . \qquad (29)$$

Combining equations (29) and (27) and then (25) we get

$$\begin{aligned}
\left(\Upsilon^{\mathrm{SOS}}(G)\right)^2 &= O\left(\frac{d\cdot\log s_{\max}}{(1-\sigma)^3}\right) + O\left(\frac{d\cdot\log^2 s_{\max}}{(1-\sigma)^3}\right) \\
&= O\left(\frac{d\cdot\log^2 s_{\max}}{(1-\sigma)^3}\right) \ .
\end{aligned}$$

Observe that

$$1-\sigma = (1-\sqrt{\sigma})(1+\sqrt{\sigma}) \ge (1-\sqrt{\sigma}) = 1-\sqrt{\beta-1} \ ,$$

where we have

$$\begin{aligned}
1-\sqrt{\beta-1} &= 1 - \frac{\lambda}{1+\sqrt{1-\lambda^2}} \\
&= \frac{1-\lambda+\sqrt{1-\lambda^2}}{1+\sqrt{1-\lambda^2}} \\
&\ge \frac{1}{2}\cdot\left(1-\lambda+\sqrt{1-\lambda^2}\right) \\
&\ge \frac{1}{2}\cdot\sqrt{1-\lambda}\cdot\left(\sqrt{1-\lambda}+\sqrt{1+\lambda}\right) \\
&\ge \frac{1}{2}\cdot\sqrt{1-\lambda} \ .
\end{aligned}$$

Therefore,

$$\Upsilon^{\mathrm{SOS}}(G) = O\left(\frac{\sqrt{d}\cdot\log s_{\max}}{(1-\lambda)^{3/4}}\right) \ .$$

This finishes the proof of the first statement. The bound in the second statement follows immediately from statement (1) and Theorem 3. ∎

V. Negative Load for SOS

In second order diffusion nodes might not have enough load to satisfy all their neighbors' demand. This situation, which we refer to as *negative load*, motivates studying by how much a node's load may become negative. Here we study the minimum amount of load that nodes need in order to prevent this event. In the following we calculate a bound on the minimum load of every node that holds during the whole balancing process. Note that, if every processor has such a minimum load at the beginning of the balancing process, there will be no processor with negative load. Hence, these bounds can also be regarded as bounds on the minimum load of every processor in order to avoid negative load.

Let $\bar{x} = (\bar{x}_1, \ldots, \bar{x}_n)$ be the balanced load vector. Define $\Delta(t) = \|x(t) - \bar{x}\|_\infty$, and $\Phi(t) = \|x(t) - \bar{x}\|_2$, where $\|.\|$ is the norm operator. Then the following observation estimates the load at the end of every step.

Observation 5. *In continuous SOS with $\beta = \beta_{opt}$ we have*

$$x(t) \geq -\sqrt{n}\cdot\Delta(0) \ .$$

Proof: We first note that

$$\Delta(t) \leq \Phi(t) \overset{(*)}{\leq} \lambda^t\cdot\Phi(0) \leq \lambda^t\cdot\sqrt{n}\cdot\Delta(0) \ , \qquad (30)$$

where $(*)$ follows from a result by Mutukrishnan et al. [19]. They show that $x(t) = M(t)\,x(0)$ for an $n\times n$ matrix $M(t)$ defined recursively. They also show that $\Phi(t) \leq \gamma(t)\cdot\Phi(0)$ where $\gamma(t)$ is the second largest eigenvalue in magnitude of $M(t)$ and $\gamma(t) \leq \lambda^t$ [19, Proof of Theorem 2]. Though they only consider homogeneous networks, their argument also applies to the heterogeneous case. The proof now follows by considering the facts $-x(t) \leq \Delta(t)$ and $\lambda^t < 1$. ∎

It should be noted that the load during a single balancing step can be lower than the bound given in in Observation 5 since Observation 5 considers only snapshots of the network at the end of each round. It might be possible that a node has to send more load items to some of its neighbors than it has at the beginning of round t, but still its load remains positive at the end of round t. This can happen if it also receives many load items from other neighbors in round t. To study the negative load issue it is helpful to divide every round in two distinct steps, where in the first step all nodes send out their outgoing flows. In the second step, they receive incoming flows sent by their neighbors in the first step. At the end of the first step all the outgoing flows are sent out but no incoming flow is yet received. To prevent negative load the load of every node has to be non-negative at this point. We call this state the *transient state* and use $\breve{x}_i(t)$ to denote the load in the transient state. Note that we always have $\breve{x}_i(t) \leq x_i(t)$ and $\breve{x}_i(t) \leq x_i(t+1)$. The following theorem provides a lower bound on $\breve{x}_i(t)$.

Theorem 10. *In a continuous SOS process with $\beta = \beta_{opt}$ we have $\breve{x}_i(t) \geq -O\left(\sqrt{n}\cdot\Delta(0)/\sqrt{1-\lambda}\right)$.*

Proof: Observe that for $t > 1$ and an arbitrary node i

$$y_{i,j}(t) = (\beta-1)\cdot y_{i,j}(t-1) + \beta\cdot\alpha_{i,j}\cdot\left(\frac{x_i(t)}{s_i} - \frac{x_j(t)}{s_j}\right) \qquad (31)$$

and since $\bar{x}_i/s_i = \bar{x}_j/s_j$ we have

$$\begin{aligned}
&\sum_{j\in N(i)} |y_{i,j}(t)| \\
&\quad\leq (\beta-1)\cdot\sum_{j\in N(i)} |y_{i,j}(t-1)| \\
&\qquad + \beta\cdot\sum_{j\in N(i)} \alpha_{i,j}\cdot\left|\frac{x_i(t)}{s_i} - \frac{x_j(t)}{s_j}\right| \\
&\quad\leq (\beta-1)\cdot\sum_{j\in N(i)} |y_{i,j}(t-1)| \\
&\qquad + \beta\cdot\sum_{j\in N(i)} \alpha_{i,j}\cdot\left(\left|\frac{x_i(t)}{s_i} - \frac{\bar{x}_i}{s_i}\right| + \left|\frac{x_j(t)}{s_j} - \frac{\bar{x}_j}{s_j}\right|\right) \ .
\end{aligned}$$

Let $g(t) = \sum_{j\in N(i)} |y_{i,j}(t)|$. Recall that $\beta < 2$, and for all i, $s_i \geq 1$ and $\sum_{j\in N(i)\cup\{i\}} \alpha_{i,j} = 1$. So we get

$$\begin{aligned}
&g(t+1) \\
&\quad\leq (\beta-1)\cdot g(t) \\
&\qquad + 2\sum_{j\in N(i)} \alpha_{i,j}\left(|x_i(t+1) - \bar{x}_i| + |x_j(t+1) - \bar{x}_j|\right) \\
&\quad\leq (\beta-1)\cdot g(t) + 4\cdot\Delta(t+1)\cdot\sum_{j\in N(i)} \alpha_{i,j} \\
&\quad\leq (\beta-1)\cdot g(t) + 4\cdot\Delta(t+1)
\end{aligned}$$

and by (30)

$$\leq (\beta-1)\cdot g(t) + 4\cdot\lambda^{t+1}\cdot\sqrt{n}\cdot\Delta(0) \ . \qquad (32)$$

Note that $g(0) < \Delta(0)$. From the recurrence of (32) we obtain

$$\begin{aligned}
g(t+1) &\leq 4\sum_{i=0}^{t}(\beta-1)^{t-i}\cdot\lambda^i\cdot\sqrt{n}\cdot\Delta(0) \\
&= 4\sqrt{n}\cdot\Delta(0)\cdot\frac{\lambda^{t+1} - (\beta-1)^{t+1}}{\lambda-(\beta-1)} \\
&\leq 4\sqrt{n}\cdot\Delta(0)\cdot\frac{\lambda}{\lambda-(\beta-1)} \qquad (33)
\end{aligned}$$

where the last inequality holds because $\lambda < 1$. On the other hand, we have

$$\begin{aligned}
\lambda - (\beta-1) &= \frac{(1+\lambda)\sqrt{1-\lambda^2} - (1-\lambda)}{1+\sqrt{1-\lambda^2}} \\
&> \frac{\sqrt{1-\lambda}\cdot(\sqrt{(1+\lambda)^3} - \sqrt{1-\lambda})}{2} \\
&> \frac{\sqrt{1-\lambda}\cdot(1-\sqrt{1-\lambda})}{2} \\
&> \frac{\sqrt{1-\lambda}\cdot\lambda}{4} \ .
\end{aligned}$$

Therefore, we can apply the above to (33) to get the bound $g(t) = O\left(\sqrt{n} \cdot \Delta(0)/\sqrt{1-\lambda}\right)$.

To complete the proof, we note that $\breve{x}_i(t) \geq x_i(t) - g(t)$, while by Observation 5 we have $x_i(t) \geq -\sqrt{n} \cdot \Delta(0)$. This yields the lower bound of $-O\left(\sqrt{n} \cdot \Delta(0)/\sqrt{1-\lambda}\right)$. ∎

The next result shows that the asymptotic lower bound obtained in Observation 5 also holds for the randomized discrete second-order process R = R(SOS) in many cases, for instance, when $s_{\max}$ is polynomial in n and $d/(1-\lambda)^{3/4} = O(n^{0.5-\varepsilon})$ for some $\varepsilon > 0$. This is true, e.g., for tori with four or more dimensions, hypercubes, and expanders. Then we can apply a similar argument as in the proof of Theorem 10 to get a lower bound for R.

Theorem 11. *In a discrete SOS process $R = R(\mathrm{SOS})$ with $\beta = \beta_{opt}$, $s_{\max}$ polynomial in n, and $d/(1-\lambda)^{3/4} = O(n^{0.5-\epsilon})$ for some $\epsilon > 0$, we have*

$$\breve{x}_i^R(t) \geq -O\left(\frac{\sqrt{n} \cdot \Delta(0) + d^2}{\sqrt{1-\lambda}}\right) \ .$$

Proof: To show this result we first rewrite (3) as follows.

$$y_{i,j}(t) \leq (\beta - 1) \cdot y_{i,j}(t-1) + \beta \cdot \alpha_{i,j} \cdot \left(\frac{x_i(t)}{s_i} - \frac{x_j(t)}{s_j}\right) + d \ ,$$

resulting in

$$g(t+1) \leq (\beta - 1) \cdot g(t) + 4 \cdot \lambda^{t+1} \cdot \sqrt{n} \cdot \Delta(0) + d^2 \ .$$

Then we rewrite (31) in the proof of Theorem 10 as follows.

$$y_{i,j}(t) \leq (\beta - 1) \cdot y_{i,j}(t-1) + \beta \cdot \alpha_{i,j} \cdot \left(\frac{x_i(t)}{s_i} - \frac{x_j(t)}{s_j}\right) + d$$

which gives us

$$g(t+1) \leq (\beta - 1) \cdot g(t) + 4 \cdot \lambda^{t+1} \cdot \sqrt{n} \cdot \Delta(0) + d^2 \ .$$

Proceeding with similar steps as in the proof of Theorem 10 we get the following bound for D.

$$\breve{x}_i^D(t) \geq -O\left(\frac{\sqrt{n} \cdot \Delta(0) + d^2}{\sqrt{1-\lambda}}\right) \ .$$

∎

VI. Simulation Results

In this section we present some simulation results for several balancing algorithms. We simulated discrete versions of both, first order and second order balancing schemes where we use *randomized rounding* as described in Section III-B for the discretization. Our main goal is to see under which circumstances SOS outperforms FOS.

We consider different networks which are based on various graph classes. A complete list of all graph types and parameters used for the simulation can be obtained from Table I. Our simulation tool is highly modularized and supports various load balancing schemes and rounding procedures. It can be used to simulate the load balancing process using multiple threads on a shared-memory machine. To fully utilize the capability of modern CPUs we used OpenMP to generate code that performs suitable instructions in parallel. The simulation was implemented using the C++ programming language. Our tests were conducted on an Intel Core i7 machine with 4 cores and 8 GB system memory.

If not stated otherwise we initialize our system by assigning a load of $1000 \cdot n$ to a fixed node v_0, where n is the number of nodes of the network, and the load of all other nodes is set to zero. Our data plotted in Figure 2, however, indicate that the amount of initial load does only have limited impact on the behavior of the simulation, especially once the system has converged.

We investigate the following metrics measuring the quality of the load distribution.

1) **Maximal local load difference.** This is the maximum load difference between the nodes connected by an edge. That is, the maximum local load difference for given load vectors $x(t)$ in a round t is defined as
$$\phi_{\text{local}}(x(t)) = \max_{\{u,v\} \in E} \{|x_u(t) - x_v(t)|\} \ .$$
2) **Maximum load.** This is the maximum load of any node minus the average load $\overline{x}$.
$$\phi_{\text{global}}(x(t)) = \Delta(t) = \max_{v \in V} \{x_v(t)\} - \overline{x}$$
3) **Potential based on 2-norm.** We compute the value of the potential function ϕ_t proposed by Muthukrishnan et al. [19] which is defined as
$$\phi_t = \phi(x(t)) = \sum_{v \in V} (x_v(t) - \overline{x})^2$$
In our plots we divided this potential by n.
4) **Impact of eigenvectors on load.** We initially compute the eigenvectors of the diffusion matrix and solve in each round t the the linear system $V \cdot a = x(t)$, which is defined over the orthonormal matrix of n eigenvectors V and the load vector $x(t)$. We then identify the *leading* eigenvector, i.e., the eigenvector with the largest $|a_i|$. The coefficients a_i for $i = 2, \ldots, n$ describe together with the eigenvectors the load imbalance completely [18]. Observe that the coefficient a_i in round t multiplied with the corresponding eigenvalue μ_i yields the coefficient in the following round $t+1$. Therefore, the largest coefficient governs the convergence rate in that step.
5) **Remaining imbalance.** This is the remaining imbalance of the converged system (see [13]), i.e., the number of tokens above average once this number starts to fluctuate and does not visibly improve any more. This imbalance does not occur in continuous systems and is due to the applied rounding in discrete systems.

In this first section we focus on the torus. For results w.r.t. other graph classes see Section VI-B.

A. Results for the Torus

Our main results are shown in Figure 1, where we plotted the simulation results using the second order scheme with randomized rounding in a two-dimensional torus consisting of 1000×1000 nodes and an average load of 1000. As in all following plots, the x-axis represents the number of rounds.

Table I. GRAPH TYPES AND PARAMETERS USED IN SIMULATION

Graph	Size	Parameter β
Two-Dimensional Torus	$n = 1000 \times 1000$	1.9920836447
Two-Dimensional Torus	$n = 100 \times 100$	1.9235874877
Random Graph (CM)	$n = 10^6$, $d = \lfloor \log_2 n \rfloor$	1.0651965147
Random Geometric Graph	$n = 10^4$, $r = \sqrt[4]{\log n}$	1.9554636334
Hypercube	$n = 2^{20}$	1.4026054847

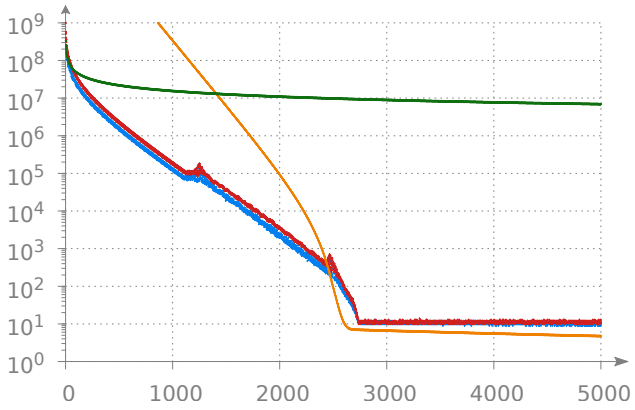

Figure 1. The maximum load minus the average load is plotted in blue, the maximum local load difference in red and the potential function ϕ_t in yellow, using SOS on a two-dimensional torus size 1000×1000. As a comparison, the green line shows the maximum load minus the average load using FOS.

The plot shows the maximum load minus the average load, the maximum local load difference, and the potential function ϕ_t on the y-axis. As a comparison, a simulation run using only first order scheme is shown as well.

It is known that the second order scheme is faster than the first order scheme w.r.t. the convergence time of the load balancing system in graphs with a suitable eigenvalue gap. However, our simulations indicate that for SOS the remaining maximal load difference does not drop below a certain threshold. Therefore, we implemented the following approach to decrease the load differences even further. First we perform a number of steps using the fast second order

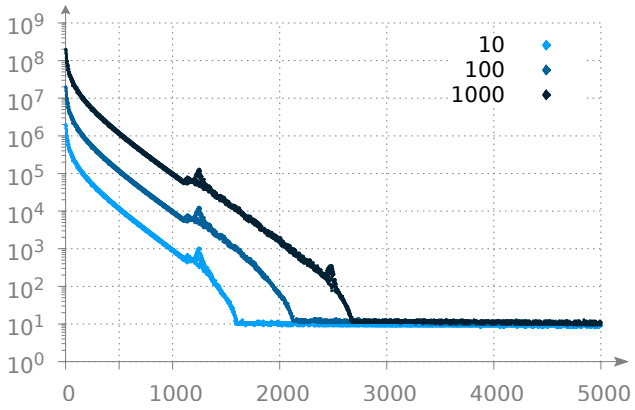

Figure 2. The plot shows the maximum load minus the average load on a two-dimensional torus of size 1000×1000. Three different initial loads were used with average loads of 10, 100, and 1000, colored from light to dark.

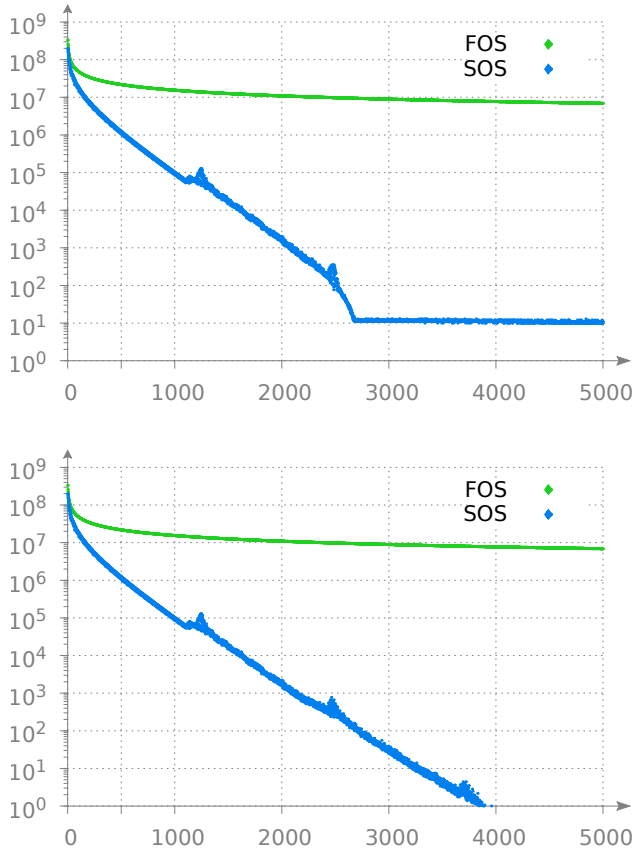

Figure 3. A comparison of the maximum load minus the average load using SOS (blue) and FOS (green) on a two-dimensional torus of size 1000×1000. The first plot shows discrete loads and randomized rounding, the second plot shows an idealized scheme.

scheme. Then, every node synchronously switches to first order scheme. We considered two different scenarios. In the first case we switched to FOS early after 2500 SOS steps. This number of steps corresponds roughly to the end of a phase of exponential decay in the potential function. In the second case we switched to FOS rather late at 3000 steps, allowing the system to run for a few hundred additional steps using the second order scheme. In both cases we observed a significant drop in both, the local and the global load differences. That is, the values for the load differences do not drop below 10 when using SOS. Once the simulation is switched to FOS, the maximum local load difference converges to a value of 4 and the maximum load minus the average load drops to 7. This is shown in Figure 4; a direct comparison is shown in Figure 5.

In the left plot in Figure 4 we furthermore observe that the load differences continue to diminish for about 200 steps (during steps 2500 to 2700) when we switch to FOS after 2500 steps. When we switch to FOS after 3000 steps (right plot in Figure 4) a drop can still be observed, however, the resulting load differences remain at a low level. To explain this behavior of the load balancing procedure we analyzed the impact of the eigenvectors of the diffusion matrix on the load balancing process. Recall that the diffusion matrix $M = (M_{ij})$

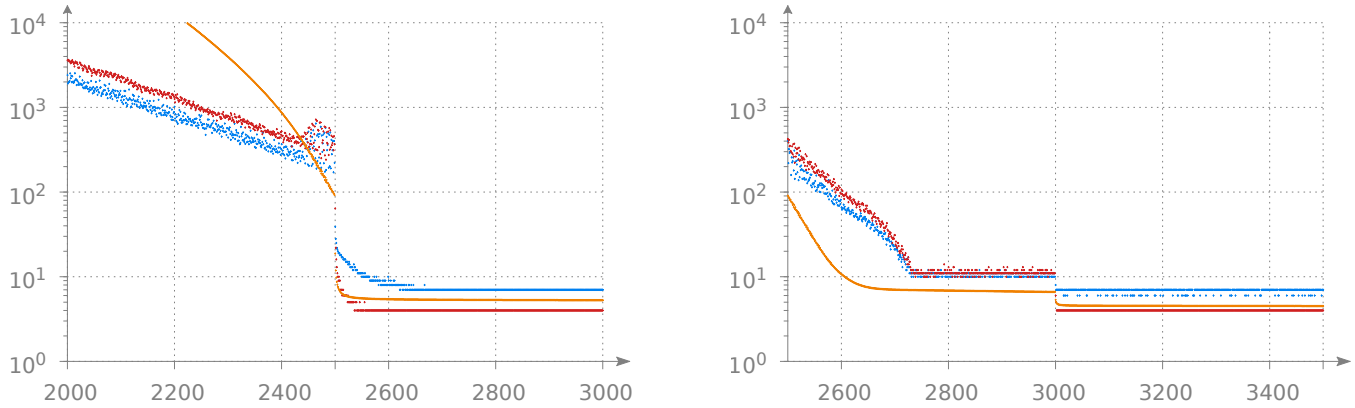

Figure 4. The plots show the maximal local difference in red, the maximal load minus the average load in blue, and the potential function ϕ_t in yellow. The simulation switches from second order scheme to first order scheme in the left and the right plot after 2500 and 3000 rounds, respectively.

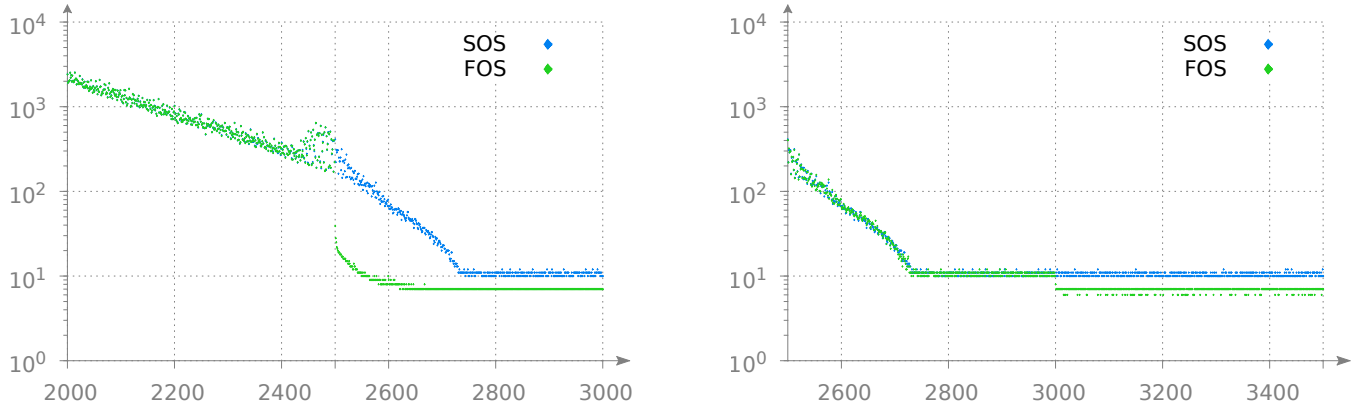

Figure 5. The plots show a direct comparison of the same data presented in Figure 4. The blue data points show the maximal load minus the average load using only a SOS approach while the green data points show the maximal load minus the average load when switching to FOS. Again, the switch has been conducted after 2500 steps in the left and 3000 steps in the right plot.

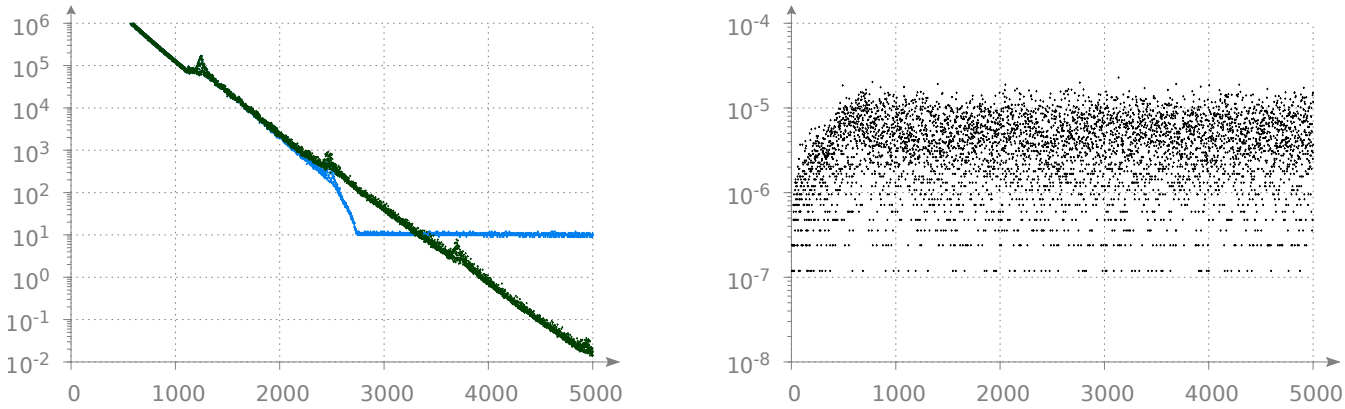

Figure 6. A comparison of the idealized second order scheme in green with a SOS using randomized rounding in blue. The idealized version is based on IEEE754 double precision floating point values as loads. The data points show the maximum load of the system minus the average load. The right plot shows the absolute value of the total load in the system at round t minus the initial total load, i.e., the absolute error.

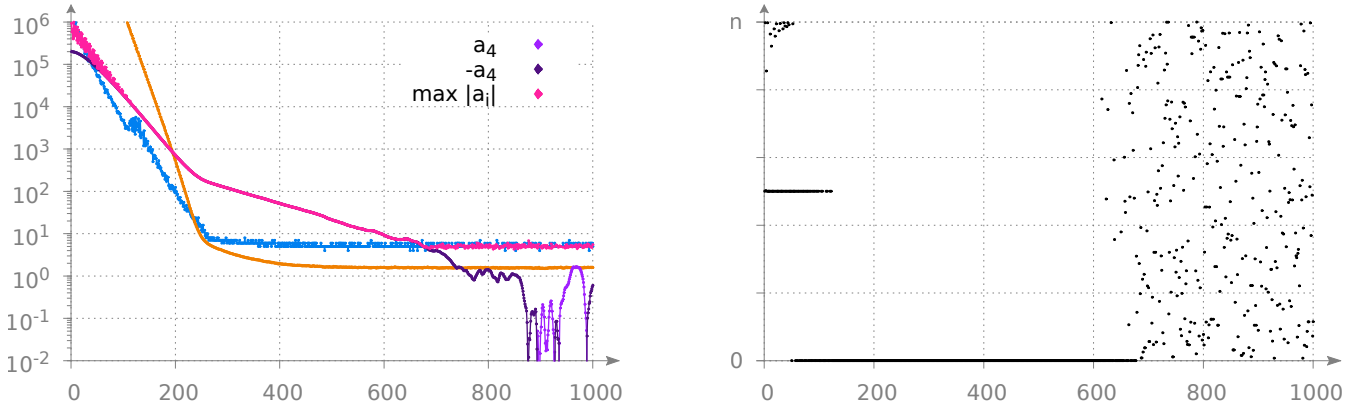

Figure 7. The left plot shows the impact of eigenvectors on the load on a two-dimensional torus of size 100×100. The maximum over all coefficients, $\max_i \{|\mathfrak{a}_i|\}$, is shown along with $\mathfrak{a}_4$. In the right plot the currently leading coefficient is shown, i.e., a black point indicates that in the given round (x-axis) the corresponding eigenvector (y-axis) has maximal impact.

is defined as

$$M_{ij} = \begin{cases} \alpha_{ij} & \text{if } i \neq j \\ 1 - \sum_{i \neq j} \alpha_{ij} & \text{if } i = j \end{cases}$$

with $\alpha_{ij} = 1/\left(\max\left\{\deg(i), \deg(j)\right\} + 1\right)$ if node i is adjacent to node j and 0 otherwise.

We used the Lapack library [4] to compute the eigenvalues and corresponding eigenvectors of M. The same library was then used to solve the set of linear systems

$$V \cdot \mathfrak{a} = W$$

for a matrix of coefficients $\mathfrak{a}$, where V denotes a matrix of eigenvectors of M and $W = (x(t))$ consists of row vectors $x(t)$ as defined in Section II containing the loads of the system at every round t. The resulting coefficients in $\mathfrak{a}$ give the impact of the corresponding eigenvectors in each round on the load. The results are shown for the torus of size 100×100 in the two plots of Figure 7. The first plot shows the maximum of these coefficients. In the simulation run corresponding to this plot we observed that starting roughly after 100 rounds this leading eigenvector corresponds to $\mathfrak{a}_4$ up until roughly round 700. After that time there is no clear leading eigenvector. This can be observed from the right plot in the same figure, where the currently leading coefficient is plotted for each round.

It seems reasonable to switch from SOS to FOS once the impact of the leading eigenvector drops below some threshold. This information, however, requires a global view on the load balancing network and therefore cannot be used in a distributed approach. In real-world applications also the trade-off between a remaining imbalance and the time required to balance the loads must be considered. We therefore investigate the effect of the time step when switching from SOS to FOS.

In Figure 8 we plotted the maximum load minus the average load for second order scheme and for an adaptive approach where we switched to FOS after a number of SOS rounds. The impact of the leading eigenvector (and the loss thereof) explains the data shown in Figure 8. Our data indicate, that once the impact of the leading eigenvector drops below a certain threshold in a round R, there is no difference in the behavior of the system when switching to FOS in some consecutive round $r \geq R$. Independently of the round R, however, we observe a significant drop in the maximum load.

Note that the maximum local load difference seems to be a good indicator for switching from SOS to FOS. Furthermore this local property is also available in a distributed system with only limited global knowledge.

In Figure 1 we also observe strong discontinuities of the local and global maximum load differences which occur approximately every 1200 to 1300 steps. To explain these discontinuities we visualized the load balancing process on the two-dimensional torus in Figure 9 as follows. We rendered a raster graphic of size 1000×1000 pixels per round. In the graphic each pixel represents a node of the torus such that neighboring pixels are connected in the network and border-pixels are connected in a periodic manner. We now set the pixels' colors to correspond to the nodes' loads, i.e., a pixel is shaded bright if its load is close to the average load and dark otherwise. In the visualization shown in Figure 9 the initial load is placed at the node with ID 0, which corresponds to the top-left pixel. Since the border-pixels *wrap around*, the loads spread in circles from all four corners, forming the *wavefronts* in the graphic. Our visualizations now indicate that the discontinuities in the local load differences and the maximum load occur whenever these wavefronts collapse at the center of the graphic, i.e., when the center node gets load for the first time. This is a consequence of the second order scheme since nodes continue to push loads towards the center pixel, even though this pixel may already have a load above average. Note that these discontinuities also occur in the idealized scheme and for smaller tori, see Figures 6 and 8, respectively.

We furthermore rendered a video of the load balancing process (available online, see [3]) which shows the behavior of the system in an intuitive way and thus helps understanding these discontinuities. Further visualizations in Figure 11 show the impact of the first order scheme. That is, after applying FOS steps the rendered image becomes more *smooth*, in contrast to the SOS steps where our visualization shows a significant amount of noise.

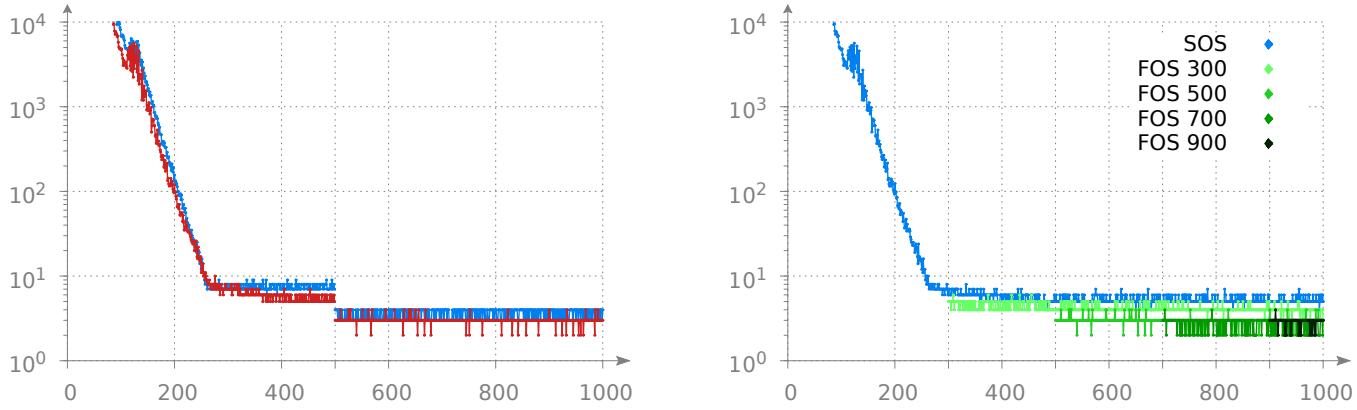

Figure 8. A plot showing the effect of switching from SOS to FOS on a two-dimensional torus of size 100×100. The left plot shows the maximum load minus the average load in blue and the maximal local load difference in red. After 500 SOS rounds the process switches to a FOS approach. In the right plot various time steps to switch from SOS to SOS are used. All data points show the current maximum load minus the average load.

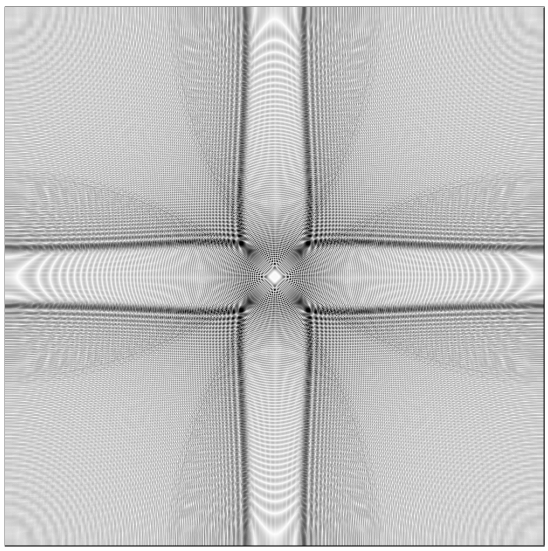

Figure 9. A visualization of the load balancing network (a two-dimensional torus of size 1000×1000) after 1100 steps. Each pixel corresponds to one node which has edges to its 4-neighborhood and is shaded such that a light pixel has a load close to the average load and a dark pixel a load close to either the maximum or minimum load of the system. Further time steps are rendered in Figure 10

To gain further insights we also implemented a simulation of the idealized load balancing procedure where loads can be split up in arbitrary small portions and any real fraction of load can be transmitted. This simulation is based on double precision floating point variables that represent the current load at a node. Therefore, a quantification takes place which introduces an error. However, we observed that in our setup the total error over all loads is small and thus can be neglected. A comparison of the idealized and discrete processes can be found in Figure 6.

B. Other Networks

For random regular graphs constructed using the configuration model [22] and the hypercube we observe only a limited improvement of SOS compared to FOS, see Figures 12 and 13, respectively. That is, the number of steps required to balance the loads up to some additive constant is only slightly larger when using FOS instead of SOS. For random graphs the remaining imbalance is the same for both FOS and SOS. For the hypercube our results indicate that the remaining imbalance using FOS is by one smaller than in the case of the SOS process. Hence, our data only show a negligible difference between FOS and SOS in these graphs. This can be related to the second largest eigenvalue of the diffusion matrix, which is $(2 + o(1))/\sqrt{d}$ for random graphs and $1 - 2/(\log n + 1)$ for hypercubes (compared to approximately $1 - \pi^2/n$ for the torus) [7]. Note that the spectral gap is also reflected in the corresponding values for β in Table I.

The random geometric graphs were generated by assigning each node a coordinate pair in the range $[0, \sqrt{n}]^2$ uniformly at random and connecting nodes v_i and v_j if and only if $d(v_i, v_j) \leq \sqrt[4]{\log n}$, where d denotes the euclidean distance. Remaining small isolated components were connected to the closest neighbor in the largest component of the graph. Even though we observe a less pronounced potential drop in random geometric graphs, the behavior of FOS and SOS in these graphs is very similar to the behavior in the torus graphs, see Figures 14 and 15.

VII. Conclusion

In this paper we analyzed a broad class of discrete diffusion type algorithms by comparing them to their continuous counterparts. Furthermore, we studied the problem of negative load in second order schemes and presented a bound for the initial minimum load in the network in order to avoid negative load during the execution of the algorithm.

Our analyses seem to provide bounds for the negative load and for the arbitrary rounding of SOS which leave room for improvement. However, in order to tighten these results, one needs some different analytic techniques. Therefore, we think that any improvement would be an interesting contribution to the field of second order diffusion schemes in particular and load balancing algorithms in general.

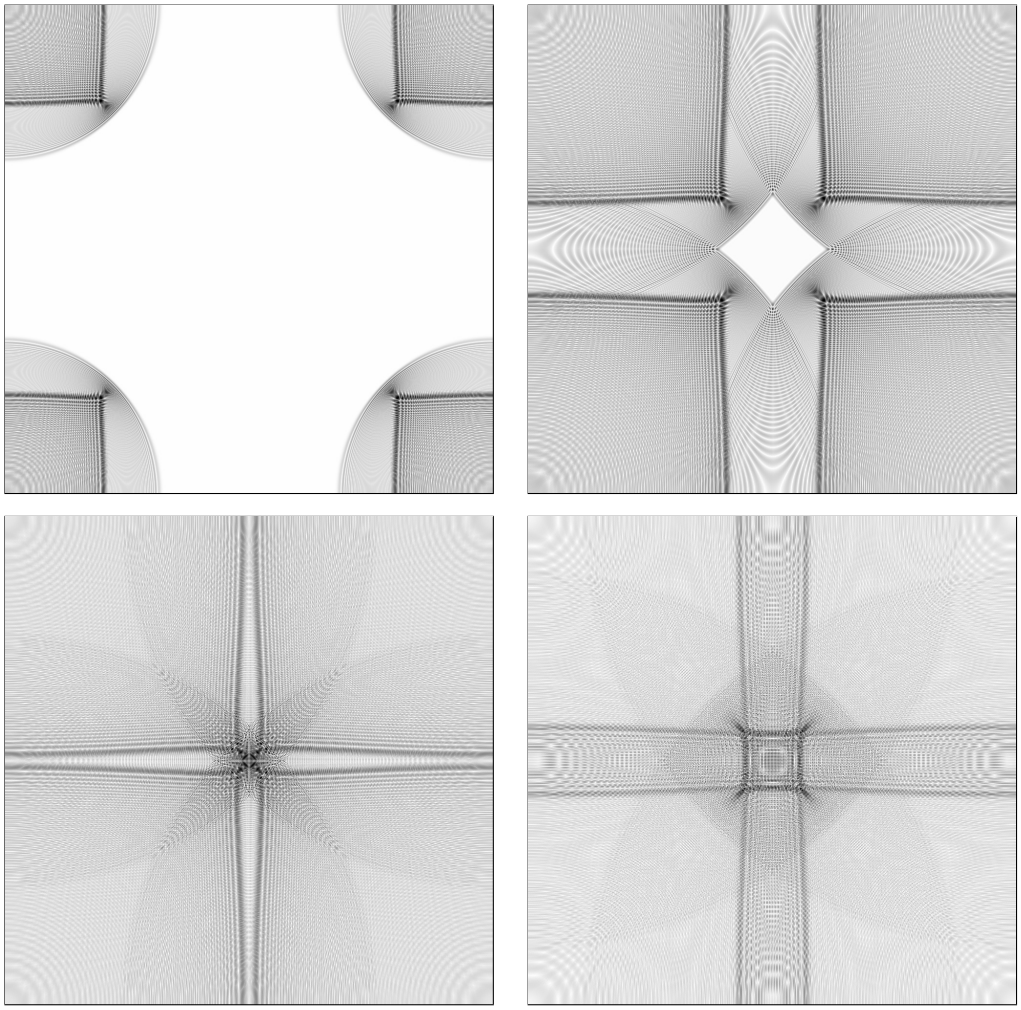

Figure 10. The figure shows the same visualization as Figure 9, rendered after 500, 1000, 1200, and 1400 steps. The network is modeled as a two-dimensional torus. Each pixel corresponds to one node which has edges to its 4-neighborhood. All pixels are shaded in an adaptive way, i.e., a light gray or white pixel indicates a load close to the average load and a dark gray or black pixel indicates a load close to either the maximum or minimum load of the system.

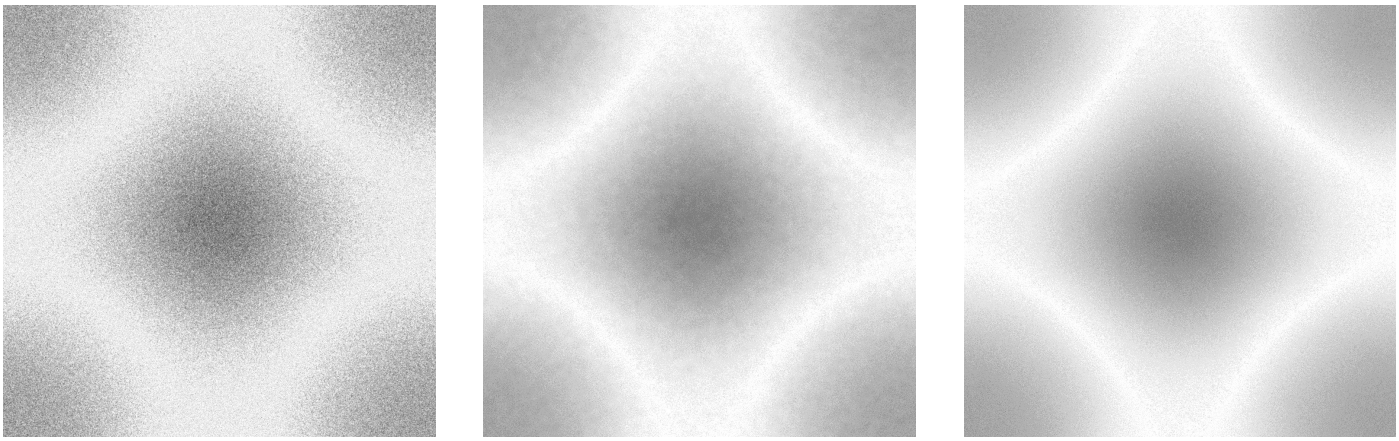

Figure 11. Above figures show the same load balancing network as Figure 9. A pixel colored white indicates a node with optimal load, a pixel colored black corresponds to a load that is more than 10 units away from the optimal value. Observe that in none of the above images such a load (which exceeds the average load by more than 10 tokens) occurs. In the center region of the left image there are several pixels which have load at least 9, whereas in the right image the maximum load exceeds the average load by at most 7. The first visualization has been rendered after 3000 SOS steps. The second image and the third image show the same network after additional 100 and 1000 FOS steps, respectively.

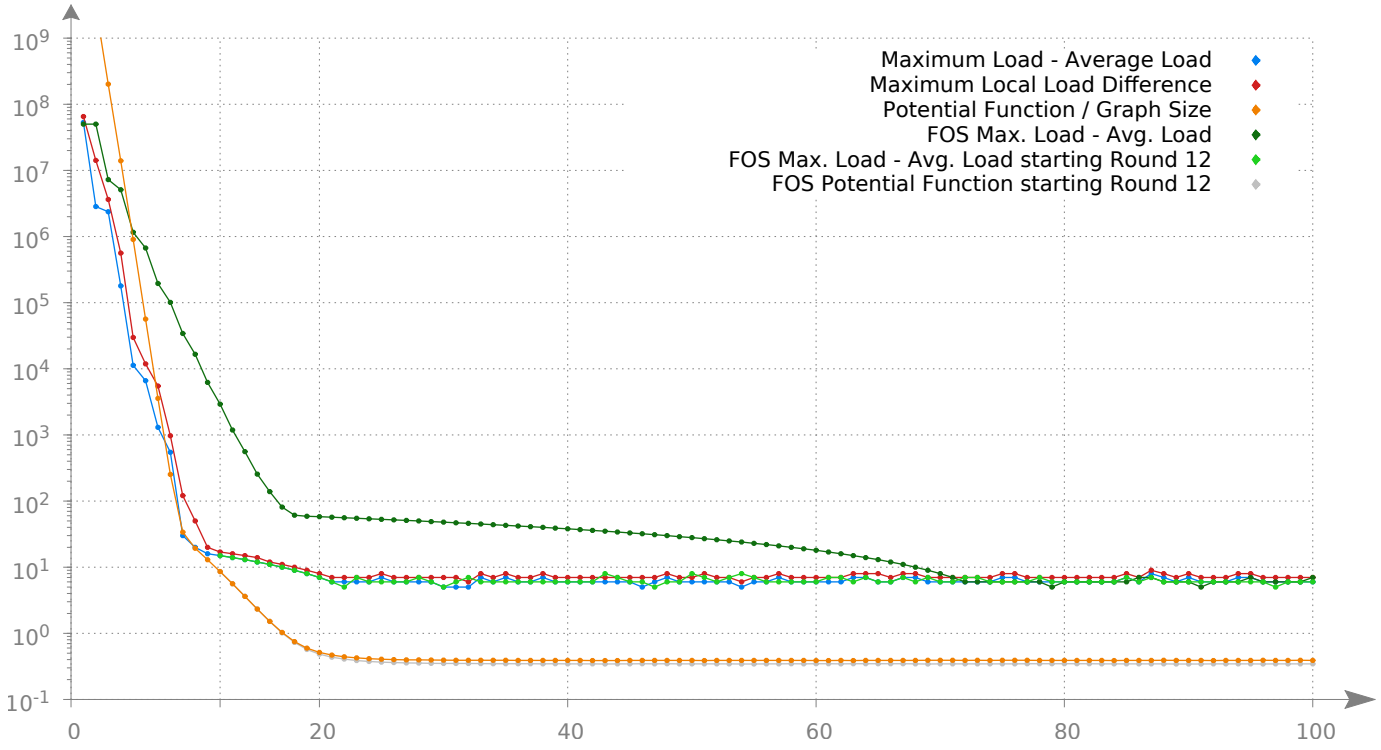

Figure 12. Load balancing simulation on a random graph in the configuration model of size 10^6 nodes with $d = 19$.

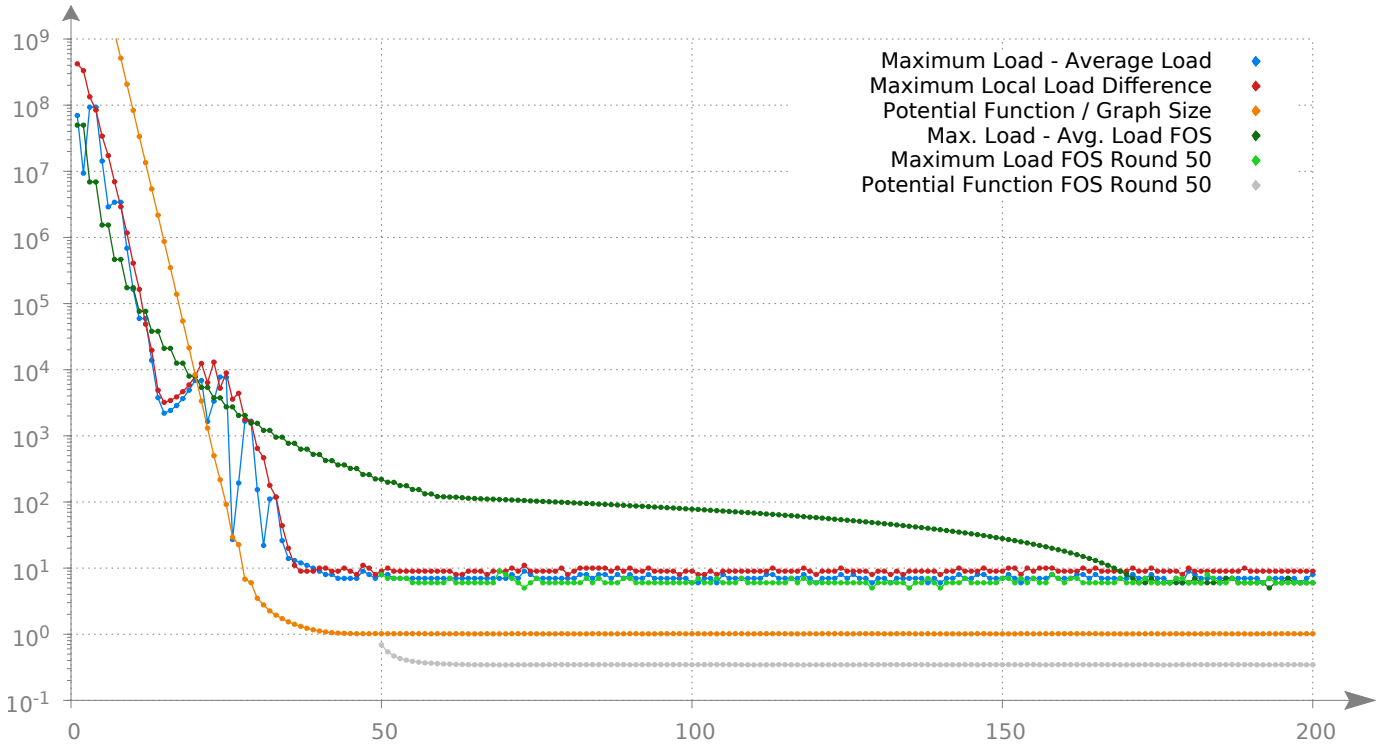

Figure 13. Load balancing simulation on a hypercube with $n = 2^{20}$ nodes. The green data points show the effect of switching to FOS after 32 steps.

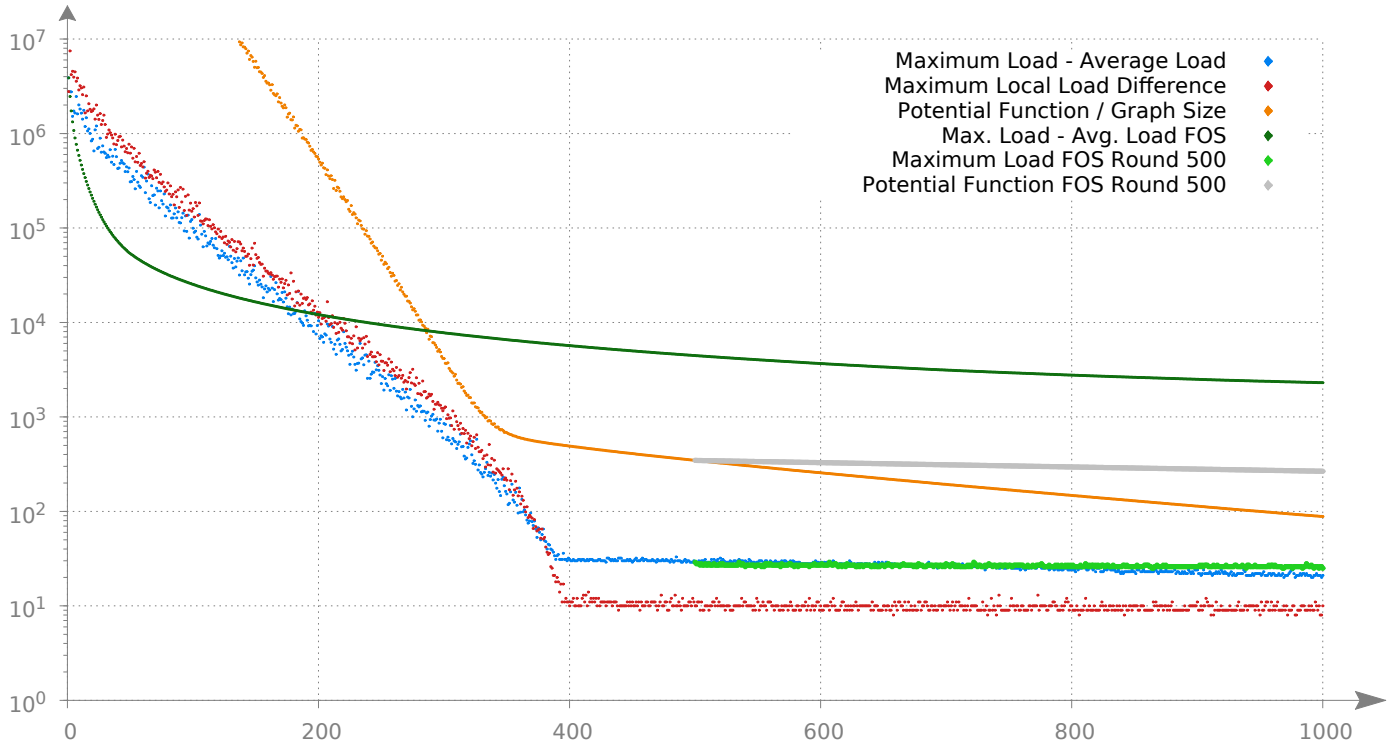

Figure 14. Load balancing simulation on a random geometric graph with 10.000 nodes in $[0, \sqrt{n}]^2$ with connectivity radius $\sqrt{\log n}$.

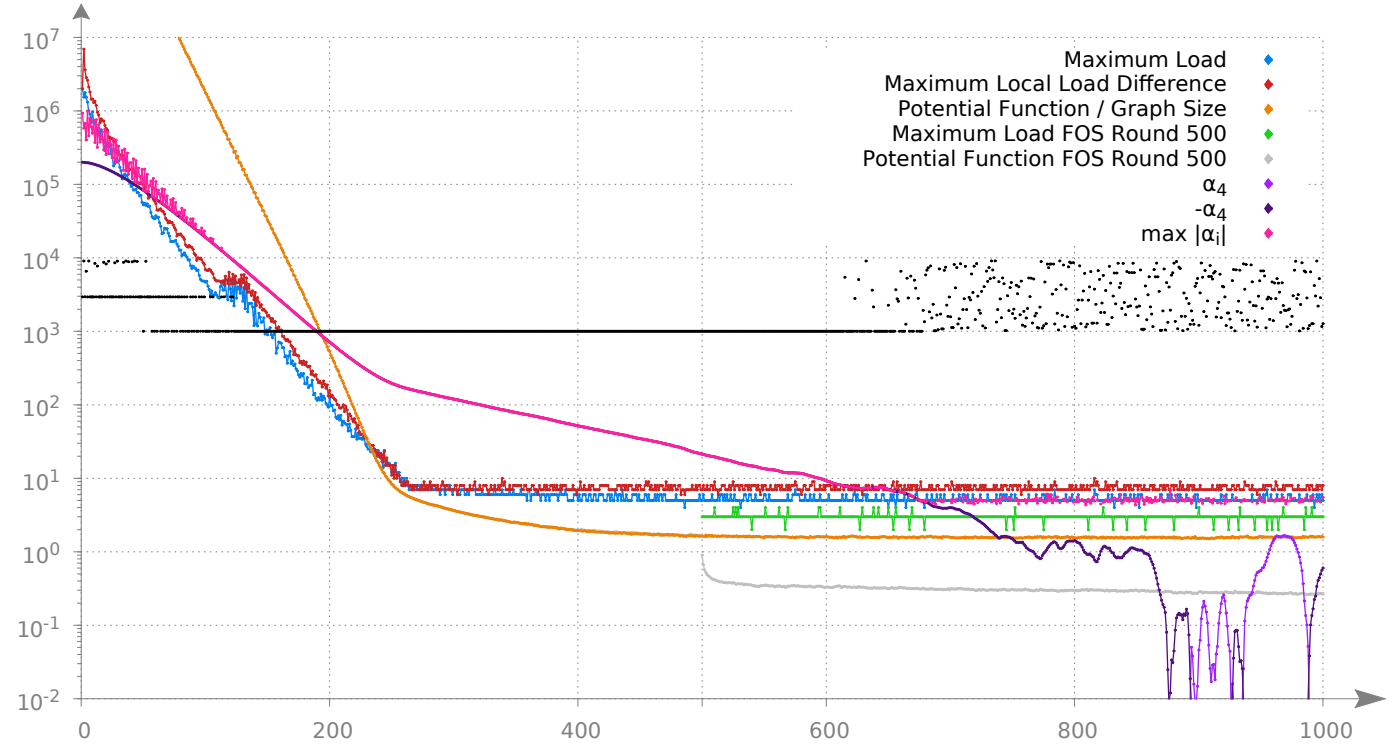

Figure 15. Load balancing simulation on a two-dimensional torus of size 100×100. The purple line shows the maximum coefficient $\max\{|\alpha_i|\}$ for the impact of the eigenvectors on the load. This coefficient is $-\alpha_4$, starting approximately in round 100 and up to approximately round 700. The black dots also shown in this plot in the range $[10^3, 10^4]$ represent the leading coefficient, where α_1 is plotted with a value of 10^3 and α_n is plotted with a value of 10^4, with a linear scale between them. Observe that after approximately 700 rounds no single eigenvector can be identified that has a leading impact on the load.

REFERENCES

[1] C. Adolphs and P. Berenbrink, "Improved bounds for discrete diffusive load balancing." in *Proc. IPDPS'12*, 2012, pp. 820–826.

[2] H. Akbari, P. Berenbrink, and T. Sauerwald, "A simple approach for adapting continuous load balancing processes to discrete settings," in *Proc. PODC'12*, 2012, pp. 271–280.

[3] H. Akbari, P. Berenbrink, R. Elsässer, and D. Kaaser, "Load balancing visualization video," 2014, https://algorithms.cosy.sbg.ac.at/downloads/load-balancing-video.mkv.

[4] E. Anderson, Z. Bai, C. Bischof, S. Blackford, J. Demmel, J. Dongarra, J. Du Croz, A. Greenbaum, S. Hammarling, A. McKenney, and D. Sorensen, *LAPACK Users' Guide*, 3rd ed. SIAM, 1999.

[5] P. Berenbrink, C. Cooper, T. Friedetzky, T. Friedrich, and T. Sauerwald, "Randomized diffusion for indivisible loads," in *Proc. SODA'11*, 2011, pp. 429–439.

[6] P. Berenbrink, T. Friedetzky, and Z. Hu, "A new analytical method for parallel, diffusion-type load balancing," *J. Parallel Distrib. Comput.*, vol. 69, no. 1, pp. 54–61, 2009.

[7] D. Cvetkovic, M. Doob, and H. Sachs, *Spectra of graphs: Theory and application*. Academic Press, 1980.

[8] R. Diekmann, A. Frommer, and B. Monien, "Efficient schemes for nearest neighbor load balancing," *Parallel Comput.*, vol. 25, no. 7, pp. 789–812, 1999.

[9] D. Dubhashi and A. Panconesi, *Concentration of Measure for the Analysis of Randomized Algorithms*. Cambridge University Press, 2009.

[10] R. Elsässer and B. Monien, "Load balancing of unit size tokens and expansion properties of graphs," in *Proc. SPAA'03*, 2003, pp. 266–273.

[11] R. Elsässer, B. Monien, and R. Preis, "Diffusion schemes for load balancing on heterogeneous networks," *Theory Comput. Syst.*, vol. 35, no. 3, pp. 305–320, 2002.

[12] R. Elsässer, B. Monien, and S. Schamberger, "Distributing unit size workload packages in heterogeneous networks," *J. Graph Algorithms Appl.*, vol. 10, no. 1, pp. 51–68, 2006.

[13] R. Elsässer and T. Sauerwald, "Discrete load balancing is (almost) as easy as continuous load balancing," in *Proc. PODC'10*, 2010, pp. 346–354.

[14] G. C. Fox, R. D. Williams, and P. C. Messina, *Parallel Computing Works!* Morgan Kaufmann, 1994.

[15] T. Friedrich, M. Gairing, and T. Sauerwald, "Quasirandom load balancing," *SIAM J. Comput.*, vol. 41, no. 4, pp. 747–771, 2012.

[16] T. Friedrich and T. Sauerwald, "Near-perfect load balancing by randomized rounding," in *Proc. STOC'09*, 2009, pp. 121–130.

[17] B. Ghosh and S. Muthukrishnan, "Dynamic load balancing by random matchings," *J. Comput. Syst. Sci.*, vol. 53, pp. 357–370, 1996.

[18] G. H. Golub and R. S. Varga, "Chebyshev semi-iterative methods, successive overrelaxation iterative methods, and second order richardson iterative methods," *Numerische Mathematik*, vol. 3, no. 1, pp. 147–156, 1961.

[19] S. Muthukrishnan, B. Ghosh, and M. H. Schultz, "First- and second-order diffusive methods for rapid, coarse, distributed load balancing," *Theory Comput. Syst.*, vol. 31, no. 4, pp. 331–354, 1998.

[20] Y. Rabani, A. Sinclair, and R. Wanka, "Local divergence of markov chains and the analysis of iterative load-balancing schemes," in *Proc. FOCS'98*, 1998, pp. 694–703.

[21] T. Sauerwald and H. Sun, "Tight bounds for randomized load balancing on arbitrary network topologies," in *Proc. FOCS'12*, 2012, pp. 341–350.

[22] N. C. Wormald, "Models of Random Regular Graphs," in *Surveys in Combinatorics, 1999*, ser. London Mathematical Society Lecture Note Series. Cambridge University Press, 1999, vol. 267, pp. 239–298.